



INDIANA UNIVERSITY

OBSERVATORY ON SOCIAL MEDIA

ANNUAL REPORT 2022-2023

OSoMe At A Glance

The Observatory on Social Media (OSoMe, pronounced *awe•some*) is an interdisciplinary research center at Indiana University studying the role of media and technology in society. Our fourth year brought us several new opportunities.

The research faculty of OSoMe led dozens of projects and supervised 16 students. **Betsi Grabe** continued our collaboration with the New Georgia Project investigating vaccine hesitancy among Black Americans and [the role Black Twitter might play in mitigating hesitancy](#) for members of this community. **Mike Gruszczynski** studied the interaction of media coverage of sociopolitical issues and citizen engagement with that coverage through social media and other online fora. **Jim Shanahan** published a major [meta-analysis](#) on the relation of social media use to beliefs and attitudes. **John Paolillo** analyzed social media, ethics, and political crisis in Sri Lanka. **Elaine Monaghan** coordinated our teaching efforts. **Filippo Menczer** and **Alessandro Flammini** led research on mining and modeling social media platforms and influence campaigns. We also hosted two visiting scholars. **Indra Bock** (Bielefeld Graduate School) joined our research efforts in the detection of bot-like activities in political conversations on Twitter. **Laura Jahn** (University of Copenhagen) worked with our teams on friction strategies to curb the production of misinformation and harmful online content.

Here we report on the major highlights of our 2022-2023 activities. Full lists of papers, public tools, media coverage, data, open-source code, and other information are available at osome.iu.edu.

2022-2023 by the numbers:

15 million \$ in submitted grant proposals

28 papers


79 students enrolled in our first course on "(mis)information literacy"

80 media hits

50 papers presented at summer 2023 conferences

OSoMe Awesome Speakers

This fall, we launched Awesome OSoMe Speakers. Led by PhD student Matthew DeVerna and our executive director Caitlin Watkins, the series is designed as a platform to showcase groundbreaking work by leading researchers on the subjects of social media manipulation, information integrity, impacts, and interventions.



Students, in particular, stand to benefit significantly from this unique opportunity to engage with experts hailing from fields such as political science, computer science, communications, and internet studies. By attending these talks, students gain a nuanced understanding of the multifaceted challenges posed by social media, fostering cross-disciplinary perspectives that help them navigate the intersections of technology, society, and information.

In the Fall 2023, **Joshua Tucker** from NYU spoke about exposure to the Russian Internet Research Agency foreign influence campaign on Twitter in the 2016 US election and its relationship to attitudes and voting behavior. **Gianluca Stringhini** from Boston University presented computational methods to measure and mitigate online disinformation. **Luca Luceri** from USC reported on AI-driven approaches for countering influence campaigns. **Franziska Roesner** from the University of Washington introduced a software tool to support effective peer response to misinformation on social media. **Brendan Nyhan** from Dartmouth College explained that subscriptions and external links help drive resentful users to alternative



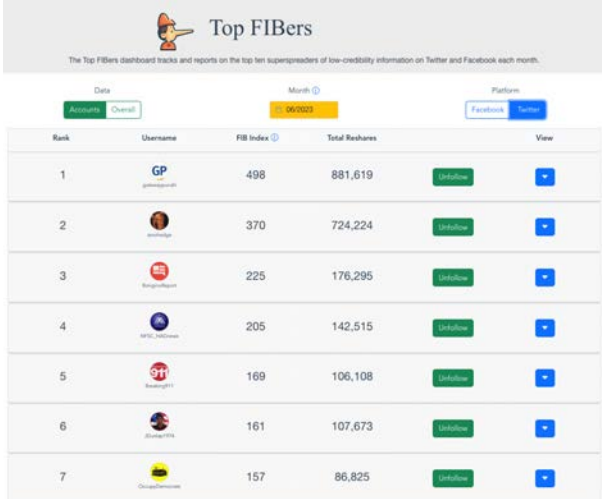
and extremist YouTube videos. **Ceren Budak** from the University of Michigan presented results on an experiment that explored the dynamics of unfollowing misinformation spreaders. And **Renée DiResta** from Stanford University spoke about generative AI and the challenge of unreality, particularly in the realm of child safety.

This spring will include talks by **David Rand** (MIT), **Sandra González-Bailón** (University of Pennsylvania), **Andrew Guess** (Princeton), **David Lazer** (Northeastern), **Kate Starbird** (University Washington), **David Broniatowski** (George Washington University), and **Sinan Aral** (MIT). For more details about upcoming awesome talks, or to watch recordings of previous talks, please visit <https://osome.iu.edu/events/speaker-series>.

Top FIBers dashboard

By rapidly disseminating information regardless of its veracity, social media superspreaders can influence consequential conversations — for better or worse — related to elections, public health and social issues.

In May 2023 we launched the new Top FIBers dashboard (osome.iu.edu/tools/topfibers), which tracks superspreaders who disseminate large quantities of low-credibility content.



The screenshot shows the 'Top FIBers' dashboard interface. At the top, there's a title 'Top FIBers' with a small icon of a person with a megaphone. Below the title, a subtitle reads 'The Top FIBers dashboard tracks and reports on the top ten superspreaders of low-credibility information on Twitter and Facebook each month.' There are three filter buttons: 'Platform' (with 'Facebook' and 'Twitter' options), 'Month' (set to '05/2023'), and 'Date' (with 'Previous' and 'Current' options). The main content is a table with columns: Rank, Username, FIB Index, Total Reshares, and View. The table lists the top 7 superspreaders.

Rank	Username	FIB Index	Total Reshares	View
1	GP	498	881,619	[View]
2	[Profile]	370	724,224	[View]
3	[Profile]	225	176,295	[View]
4	[Profile]	205	142,515	[View]
5	[Profile]	169	106,108	[View]
6	[Profile]	161	107,673	[View]
7	[Profile]	157	86,825	[View]

This dashboard provides monthly reports highlighting the top ten superspreaders of low-credibility information on social media. Superspreaders are identified using a metric called the “False Information Broadcaster index” or “FIB index,” which captures the consistency with which users share links to low-credibility sources that are subsequently reshared many times (preprint [doi:10.48550/ARXIV.2207.09524](https://doi.org/10.48550/ARXIV.2207.09524)).

For example, a user with a FIB index of 100 has shared at least 100 posts linking to low-credibility sources, each of which has been reshared at least 100 times. On the other hand, a user who posts only one post linking to a low-credibility source will have a FIB index of one, even if it was reshared millions of times.

“Research from our Observatory and others has shown that a few influencers are responsible for a large proportion of low-credibility content being shared online, including harmful content such as false vaccine claims,” says Matt DeVerna, the OSoMe PhD student who has coordinated the project. “Our new dashboard will help citizens understand the role of these bad actors in the spread of misinformation.”

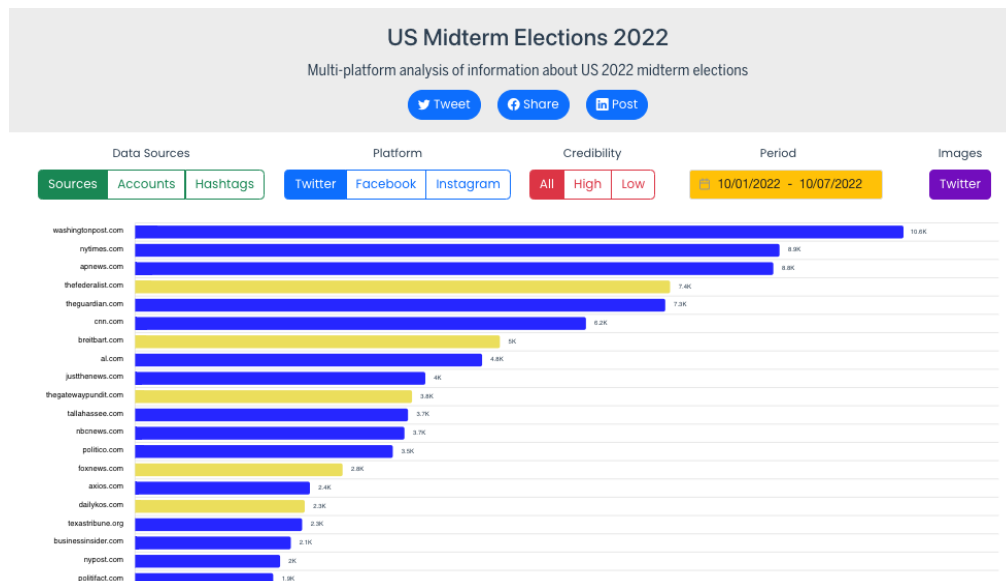
OSoMe uses a definition of misinformation commonly utilized in academic research, which focuses on sources of information that mimic news media content in form but not in organizational process or intent. Using this definition, the dashboard searches for posts that contain at least one link to any source that meets three criteria: it is rated to have low credibility, it is categorized as either “conspiracy/pseudoscience” or “questionable/fake news,” and it is labeled as having a low or very low factual score. These ratings are curated by an independent third party, Media Bias Fact Check, and compiled by Iffy.news.

“As the social media landscape evolves, OSoMe is looking for new ways to partner with platforms on making data accessible to researchers,” said Caitlin Watkins, OSoMe’s executive director. “This tool exemplifies the public good that comes from such data access.”

Tool and Data to Monitor Midterm Election Discourse on Social Media

Social media are utilized by millions of citizens to discuss important political issues. Politicians use these platforms to connect with the public and broadcast policy positions. Therefore, data from social media has enabled many studies of political discussion. While most analyses are limited to data from individual platforms, people are embedded in a larger information ecosystem spanning multiple social networks.

In the lead up to the 2022 US Midterm election, OSoMe began collecting social media posts from Twitter, Facebook, Instagram, Reddit, and 4chan. The collection contains links to posts about the midterm elections based on a comprehensive list of keywords and tracks the social media accounts of 1,011 candidates from October to December 2022. The Indiana University 2022 U.S. Midterms Multi-Platform Social Media Dataset (MEIU22) is publicly available to researchers (Proc. ICWSM 2023, [doi:10.1609/icwsm.v17i1.22205](https://doi.org/10.1609/icwsm.v17i1.22205)).



To provide up-to-date insights from the posts the team collected, we developed a dashboard designed to provide insight into the spread of information during the election cycle. The US Midterm Elections 2022 dashboard (osome.iu.edu/tools/midterm22) highlights the most widely shared websites and the accounts responsible for spreading this content across three major social media platforms: Twitter, Facebook, and Instagram. Users can also access information on the top hashtags and images circulating on these platforms.

Taking a nuanced approach to credibility, the dashboard categorizes links as either "high-" or "low-credibility" based on ratings provided by NewsGuard. Additionally, the tool highlights the accounts responsible for sharing the most widely circulated articles, offering a granular view of information sources.

Funded Projects



Our collaboration continues with the University of Southern California and the University of Maryland, supported by the **DARPA** project titled “Universal Population Response Characterization Algorithms for OnLine Environments.” Our team is engaged in the development of machine learning methods tailored for the identification and characterization of both the users involved in promoting online campaigns and the groups susceptible to their influence.


“Bots Building Bridges (3B): Theoretical, Empirical, and Technological Foundations for Systems that Monitor and Support Political Deliberation Online” is a collaboration between the University of Bielefeld, the University College Dublin, and Indiana University, supported by a grant from the **Volkswagen Foundation**. The overarching objective is to delve into the intricate social dynamics between technical systems, human actors, and partisan agendas within the online environment. Our team has been developing AI algorithms to detect and study inauthentic accounts and automated behaviors online.



Working with partners at the University of Applied Sciences and Arts of Southern Switzerland, the University of Zurich, and the University of Southern California, our team received an award from the **Swiss National Science Foundation** to support a “Call for

Regulation Support In Social Media” (CARISMA). Inputs from media policy and governance research will be used to formulate a set of policy alternatives to mitigate online harm, while computational social science models will evaluate the effects and impacts of moderation policies. CARISMA will recommend effective and transparent policy interventions that platforms and regulators can use to mitigate the harms of social media actors responsible for abusive and illicit behaviors.

The **Open Technology Fund** supports our development of several online tools to study social media dynamics and manipulation. The initial objective of this project was to enhance the accessibility, robustness, and feature richness of OSoMe tools, empowering communities to assess the reliability of information and sources independently. The discontinuation of data availability from Twitter in mid-2023 forced our team to make some changes. In collaboration with OTF staff, we revised our project scope to ensure the continued viability of our tools that provide valuable resources for the broader online community. We are also developing a new tool that will detect and visualize coordinated influence campaigns on platforms beyond Twitter, possibly including Mastodon, BlueSky, TikTok, and others.



Supported by the **Alfred P. Sloan Foundation** through the Social Science Research Council's **Mercury Project**, a new collaborative effort with the University of Maryland, Dartmouth College, and Tel Aviv University is dedicated to investigating strategies for mitigating the impact of low-credibility information sources on social media. In particular, our interdisciplinary team is examining the long-term effects of altering information environments by encouraging users to mute untrustworthy accounts.

Continuous funding from the **Knight Foundation** has been instrumental for OSoMe's multidisciplinary research efforts. Their commitment has supported the OSoMe staff, maintained our data infrastructure, enhanced existing projects, and enabled us to embark on new avenues of exploration. For example, students supported by the Foundation delved into both helpful and harmful applications of AI, such as scaling up fact-checking and controlling networks of malicious social bots; developed agent-based models to study online manipulation strategies and tactics and mass-media effects; designed knowledge graph methods that could find applications in computational fact-checking; evaluated the effectiveness of vaccine safety messages targeting communities of color; and conducted experiments measuring psychophysiological responses to exposure to misinformation.



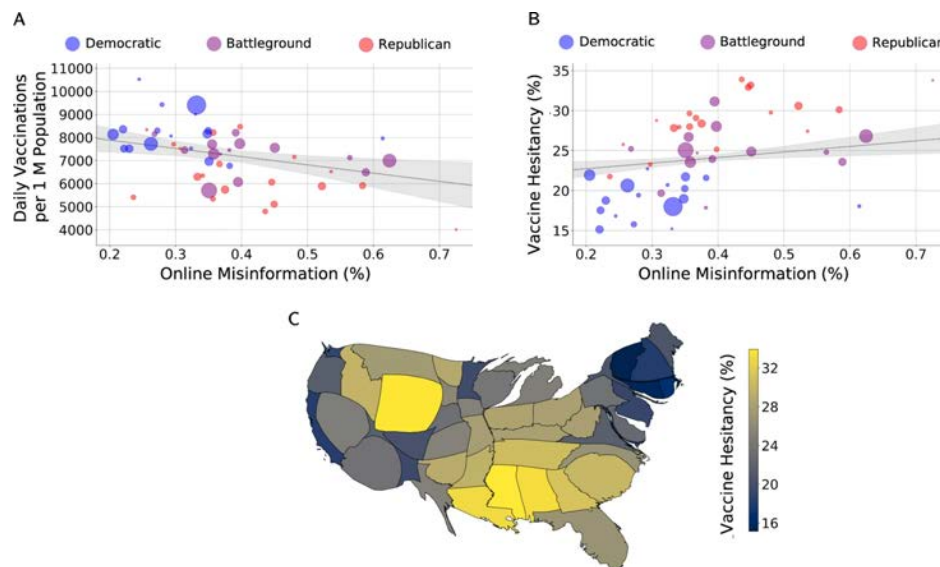
**KNIGHT
FOUNDATION**

Research Area Updates

Vaccine Hesitancy and Social Media

In early 2021, OSoMe announced the development of CoVaxxy, a dashboard visualizing the relationship between COVID-19 vaccine adoption and online (mis)information in the United States. This project became the foundation for one completed grant, two grant proposals under development, and several papers.

Our paper "Online Misinformation is Linked to Early COVID-19 Vaccination Hesitancy and Refusal" was included in the 2023 Economic Report of the U.S. President. Adding to its accolades, this paper was also recognized by *Scientific Reports* as one of the top 100 articles in 2022; it has been cited 137 times as we write this report. This paper delves into the relationship between online misinformation and the early hesitancy and refusal of COVID-19 vaccination. Widespread vaccine uptake is crucial for achieving herd immunity, but we observed variations in uptake rates across U.S. states and counties during the initial six months of the vaccination program. Our work revealed a negative relationship between these vaccination rates and online misinformation spreading in the same locations. These associations are significant even when accounting for political, demographic, and socioeconomic factors.



In one follow-up paper, we establish a causal link between real-world exposure to antivaccine content and vaccine uptake. We present a compartmental epidemic model that includes vaccination, vaccine hesitancy, and exposure to antivaccine content. We fit the model to observational data to determine that a geographical pattern of exposure to online antivaccine content across US counties is responsible for a pattern of reduced vaccine uptake in the same counties. We find that exposure to antivaccine content on Twitter

caused about 750,000 people to refuse vaccination between February and August 2021 in the US, resulting in at least 29,000 additional cases and 430 additional deaths. This work provides a methodology for linking online speech to offline epidemic outcomes. Our findings should inform social media moderation policy as well as public health interventions.

In another paper, we are exploring how misinformation can influence the spread of disease more generally. We combine a large-scale collection of social media posts discussing COVID-19 with county-level voting records and cellphone mobility data to reproduce a physical contact network consisting of nearly 20 million nodes, some marked as misinformed. We then introduce an agent-based model that accounts for the presence of misinformed individuals in the population. Exploring theoretically extreme scenarios, we estimate an upper bound for the negative effects of misinformation on the spread of the disease through the full population, estimating that as many as 47 million additional Americans could become infected at a cost of \$143 billion.

X/Twitter, Data, Bots, and New Social Media Platforms

The social media landscape has changed dramatically in recent months. After X/Twitter was acquired by Elon Musk, its data API, which had enabled so much of our research on social media use and abuse, became unavailable to researchers. Our Observatory was among the [signatories](#) of two open letters denouncing the [harm of such a decision to the research community](#) and to social media transparency. While the Twitter API was still available, we contributed to the arguments between Musk and Twitter about the prevalence of social bots [1] and how their moderation affected free speech [2]. We conducted a study on public bot perceptions [3] and published methodological papers for computational social scientists doing research on social bots [4] and on voters [5] using data from Twitter. Finally, we wrote an opinion on the promise and fallacies of community-based moderation [6]. While this kind of research is no longer possible on Twitter, new platforms are emerging. We are currently working on data collection infrastructure to support research on platforms like TikTok, Bluesky, and Mastodon.

[1] [How many bots are on Twitter? The question is difficult to answer and misses the point.](#) Menczer, F. and Yang, K., *The Conversation*

[2] [Elon Musk is wrong: research shows content rules on Twitter help preserve free speech from bots and other manipulation.](#) Menczer, F., *The Conversation*

[3] [Exposure to social bots amplifies perceptual biases and regulation propensity.](#) Yan, H. Y.; Yang, K.; Shanahan, J.; and Menczer, F. *Scientific Reports*

[4] [Botometer 101: Social bot practicum for computational social scientists.](#) Yang, K.; Ferrara, E.; and Menczer, F., *Journal of Computational Social Science*

[5] [How Twitter data sampling biases U.S. voter behavior characterizations.](#) Yang, K.; Hui, P.; and Menczer, F., *PeerJ Computer Science*

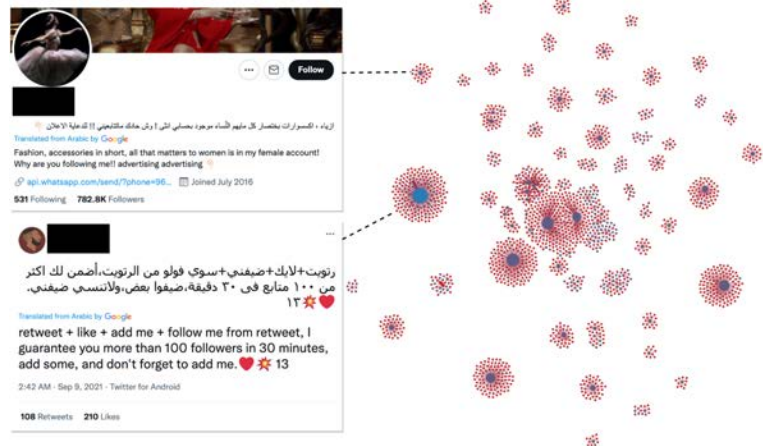
[6] [Can crowdsourcing rescue the social marketplace of ideas?.](#) Yasseri, T. and Menczer, F., *Comm. ACM*

Coordinated Online Behavior

Following up on our earlier [work about methods to detect coordinated inauthentic networks on social media](#), we published two papers on online manipulation tactics involving coordinated behaviors. One explores the use of “follow trains” to manipulate follower networks and boost partisan echo chambers [7]. Accounts heavily promoted by follow trains profit from a median

six-fold increase in daily follower growth. Some accounts using these tactics have been active for years, suggesting that platforms need to pay greater attention to this kind of abuse. Another tactic consists of flooding and deleting content [8]. Analyzing more than a billion deletions by over 11 million accounts, we uncovered coordinated networks of accounts engaged in repetitive likes and unlikes of content that is

eventually deleted, for the purpose of manipulating ranking algorithms. Finally, we introduced a language framework called BLOC for modeling social media account behaviors and capturing suspicious similarities [9]. This approach outperforms state-of-the-art methods in the detection of coordinated inauthentic behavior. We are currently working on a new online tool based on BLOC and on extensions to capture more nuanced behavioral relationships.



[7] [The Manufacture of Partisan Echo Chambers by Follow Train Abuse on Twitter.](#) Torres-Lugo, C.; Yang, K.; and Menczer, F., *In Proc. Intl. AAAI Conf. on Web and Social Media (ICWSM)*

[8] [Manipulating Twitter through Deletions.](#) Torres-Lugo, C.; Pote, M.; Nwala, A.; and Menczer, F., *In Proc. Intl. AAAI Conf. on Web and Social Media (ICWSM)*

[9] [A Language Framework for Modeling Social Media Account Behavior.](#) Nwala, A. C.; Flammini, A.; and Menczer, F., *EPJ Data Science*

Social Media Simulators

We are continuing to develop an agent-based model called SimSoM to simulate information diffusion on social media, which can be used to explore the effects of manipulation strategies and tactics [10] as well as moderation interventions [11]. With support from the CARISMA project funded by the Swiss National Science Foundation, we are enriching the model with more realistic newsfeed representations and ranking algorithms, as well as empirically-calibrated parameters. Eventually the model will be applied to evaluate the intended and unintended consequences of policy regulation.

[10] [Vulnerabilities of the Online Public Square to Manipulation. Truong, B. T.; Lou, X.; Flammini, A.; and Menczer, F., *Technical Report*](#)

[11] [Friction Interventions to Curb the Spread of Misinformation on Social Media. Jahn, L.; Rendsvig, R. K.; Flammini, A.; Menczer, F.; and Hendricks, V. F., *Technical Report*](#)

Helpful and Harmful Applications of AI

Given the striking abilities of large language models (LLMs) [12], the research community has been paying a lot of attention to helpful and harmful applications of AI. On the helpful side, we found for example that LLMs can be quite accurate in rating news credibility [13]. This inspired a randomized controlled experiment in which we tested the effect of using ChatGPT to scale up fact-checking [14]. Surprisingly, the AI did not significantly affect participants' ability to discern headline accuracy or share accurate news. However, the AI fact-checker was harmful in specific cases where its labels were inaccurate. When participants were given the option to view AI fact checks and chose to do so, they were significantly more likely to share both true and false news but only more likely to believe false news. Our findings highlight an important source of potential harm stemming from AI applications and underscore the critical need for policies to prevent or mitigate such unintended consequences. AI can also be used intentionally in harmful ways. We uncovered a Twitter botnet involving over a thousand fake accounts that employed ChatGPT to generate human-like content [15]. These accounts posted machine-generated harmful content and stolen images, and engaged with each other through replies and retweets. While the accounts in the AI botnet could be detected through their coordination patterns, current state-of-the-art LLM content classifiers fail to discriminate between them and human accounts in the wild. These findings highlight the threats posed by AI-enabled social bots. We discussed the factuality challenges of LLMs and research directions to mitigate these threats in two opinion articles [16,17].

[12] [The Inexplicable Efficacy of Language Models. Aiyappa, R. and Kachwala, Z., *XRDS*](#)

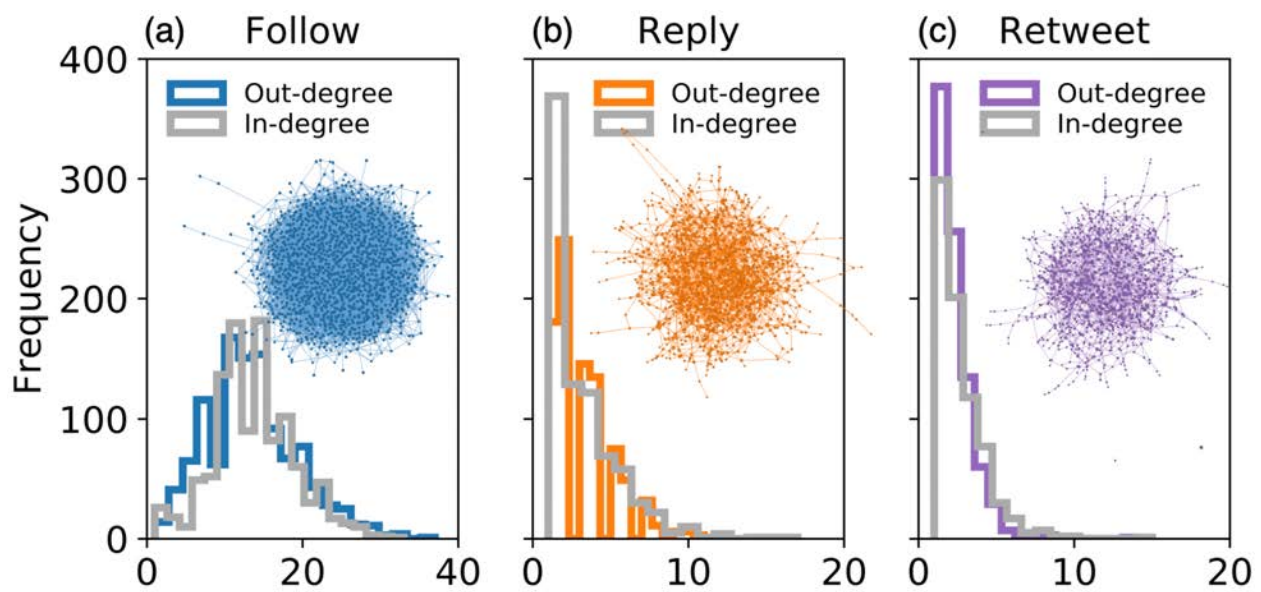
[13] [Large language models can rate news outlet credibility. Yang, K. and Menczer, F., *Technical Report*](#)

[14] [Artificial intelligence is ineffective and potentially harmful for fact checking.](#) DeVerna, M. R.; Yan, H. Y.; Yang, K.; and Menczer, F., *Technical Report*

[15] [Anatomy of an AI-powered malicious social botnet.](#) Yang, K. and Menczer, F., *Technical Report*

[16] [Factuality Challenges in the Era of Large Language Models.](#) Augenstein, I. & al., *Technical Report*

[17] [Addressing the harms of AI-generated inauthentic content.](#) Menczer, F.; Crandall, D.; Ahn, Y.; and Kapadia, A., *Nature Machine Intelligence*



Meet Our Newest Researchers!

Tuğrulcan "Tj" Elmas, Postdoctoral Fellow



Tj received his Ph.D. from EPFL, Switzerland and his B.S. from Bilkent University, Turkey, with a major in Computer Science and a minor in Philosophy. His main research focuses on computational approaches to understand and counter social media manipulation. Dr. Elmas will join the University of Edinburgh, UK as a Lecturer in Computational Social Science in Spring 2024.

Sangyeon Kim, Postdoctoral Fellow



Sangyeon holds a dual Ph.D. in Political Science and Social Data Analytics from Pennsylvania State University. He uses computational methods and experimental approaches to study internet politics, misinformation diffusion, conflict forecasting, and contentious politics, mostly in the context of the United States. His research has appeared or is forthcoming in *European Journal of Political Research*, *Political Science Research and Methods*, and *Proceedings of the International AAAI Conference on Web and Social Media*.

Filipi N. Silva, Research Scientist



Filipi joined OSoMe in July 2023 as a Research Scientist, after working as an Assistant Research Scientist for four years at the IU Network Science Institute. Filipi received his Ph.D. in Physics at the São Carlos Institute of Physics in Brazil. His primary research interests include studying and implementing new techniques for analyzing, modeling, and understanding real-world systems through a joint combination of complex networks, machine learning, and data visualization. His current research focus is on developing a new set of interactive network visualization tools for researchers.

OSoMe Students



Rachith Aiyappa is a Ph.D. candidate in Informatics with a minor in Computer Science. His research includes modeling the spread of beliefs and behavior on social networks and building pipelines using large language models to extract beliefs from social media to validate various models of belief dynamics. He is also interested in simulating artificial agents using LLMs and, in particular, measuring the (mis)alignment between the interaction among artificial agents and that among humans.



David Axelrod is a Ph.D. student in Complex Networks and Systems. Having a background in history, David is integrating research questions from history and STS with complex systems and network science approaches. Key areas of interest include information diffusion, contagions in social networks, using archival material as a data source for computational analyses, and incorporating historical analysis into examinations of technologically implicated developments.



John Bollenbacher successfully defended his thesis on December 4th! His work focuses on measuring the offline effects of online social media and he has created a method for linking online causes to offline outcomes with modern data science tools and causal inference methods. John is currently working as a Research Data Scientist at the Research Triangle Institute's Center for Data Science and AI.



Matthew DeVerna, an Informatics PhD candidate, specializes in computational social science and social media dynamics. His paper on misinformation superspreaders received a Best Student Extended Abstract Award at MISDOOM 2022. His recent work explores generative AI and fact-checking. His research, showcased at the EDMO Scientific Conference on Disinformation, explores how misinformation can amplify the spread of disease. Matthew serves as guest editor for EPJ Data Science's special issue on Computational Approaches for Cyber Social Threats.



Zoher Kachwala, a Ph.D. candidate in Computer Science, specializes in structured representations for language. His research centers on leveraging and comparing knowledge graphs derived from text for enhanced structured reasoning in NLP. His long-term goal is to use knowledge graphs to mitigate the hallucination challenges encountered by large language models in multihop reasoning tasks like question answering and fact-checking.



Munjung Kim is currently pursuing her Ph.D. in Informatics. She focuses her research on the science of science, exploring the mechanisms driving human innovation, the evolution of scientific disciplines, and the formulation of scientific theories. Her work also includes the study of hate speech and its connection to misinformation.



Kevin Mudavadi is a Ph.D. Candidate at the Media School. His research centers on the effects of online mis/disinformation on individual decision-making. Additionally, he is intrigued by the relationship between journalistic roles and media perceptions among audiences, particularly in the context of misinformation and media trust. Kevin has examined media perception, disinformation, information disorder, and journalism in the Kenyan and United States contexts.



Jimmy Ochieng is pursuing a PhD at the Media School. His research areas include political communication, media effects and political behavior, and mediation and reconciliation. He has actively contributed to peace initiatives in Africa working with the African Union Panel of Eminent African Personalities, United Nations, Centre for Humanitarian Dialogue, and Berghof Foundation.



Haley Pierce is a Ph.D. student in the Media School. Her research interests concern the political, economic, and social factors that shape individuals' political information environments and the effects of such information on people's political attitudes and behavior. In her current work, Haley is conducting experiments measuring psychophysiological responses to exposure to misinformation.



Manita Pote is a PhD student in Informatics. Her research interests encompass computational social science, network science, machine learning, and social media data mining. Manita's primary focus is on coordination detection, where she explores the identification of tactics and strategies within influence campaigns on social media. Manita received the NSF NRT Research Fellowship in 2020.




Ozgur Can Seckin is a Ph.D. student in Complex Networks and Systems. His research interests encompass computational social science, mental health concerns on social media, science of science, diffusion of false information on social media, and modeling opinion dynamics by incorporating the cognitive belief systems of individuals. Ozgur was named a Fall 2023 Luddy Grad Fellow.



Danishjeet Singh is a Computer Science undergraduate student with specialization in Artificial Intelligence and minors in Data Science and Statistics. His research interests are at the intersection of social media misinformation and computer vision. Currently, he's working on detecting and characterizing the usage of Generative AI imagery on social media platforms.



Bao Tran Truong is a Ph.D. candidate in Complex Networks and Systems. She studies ways to combat online misinformation. Her work spans from detecting inauthentic actors to developing computational models to measure the impact of social media abuse and propose moderation policies. Beyond that, she envisions a revamped online governance that highlights the role of social norms in creating a more trustworthy digital landscape.



Lake Yin is our most junior Informatics PhD student. His research interests broadly span the intersection between social networks and artificial intelligence. He is currently working on detecting coordinated influence campaigns.



Where Are They Now?

Mel Allen (they/them) is a second-year PhD student working with Dr. David Lazer at Northeastern University's Network Science Institute. They have a Bachelor's degree in math and computer science and a minor in art history from Indiana University. At OSoMe, Mel studied images posted on Twitter during the Stop Asian Hate movement and performed an analysis of political advertisements posted on Facebook regarding the ongoing conflict between Russia and Ukraine. Currently, Mel is studying how online aesthetic movements can be appropriated by extremist or far-right groups, specifically the Cottagecore to Tradwife Pipeline on Tumblr.



Harry Yan is currently a visiting assistant professor in the Department of Communication & Journalism at Texas A&M University. He obtained his dual PhD in Informatics (Complex Networks and Systems) and Media Studies at Indiana University under the joint supervision of Prof. Jim Shanahan and Prof. Filippo Menczer. His research concerns the effects of media technologies on public opinion formation and focuses on investigating how interactions with social bots affect public perceptions and policy preferences about those accounts, as well as how partisans form opinions in a high-choice multi-platform media environment.

Kai-Cheng Yang is a postdoctoral research associate at the Network Science Institute, Northeastern University. He obtained his Ph.D. in Informatics (Complex Networks and Systems) from the Luddy School of Informatics, Computing, and Engineering at Indiana University Bloomington. He is interested in computational social science with a focus on identifying bad actors, such as malicious social bots, and studying the diffusion of misinformation on social media. His current research aims to uncover how generative AI is used for deceptive and disruptive purposes and develop countermeasures leveraging AI technologies.



Staff Updates

Following a reorganization of interdisciplinary research centers at Indiana University, staff members from the IU Network Science Institute were integrated into the Observatory on Social Media. As a result, OSoMe broadened its mission to support research in social media and network science. We welcomed several new staff members.



Ben Serrette
IT Director



Sarah Beverton
Full Stack Developer



Pasan Kamburugamuwa
Full Stack Developer



Nick Liu
System Developer



Lourdes Gonzalez
Administrative Coordinator



Filipi Silva
Research Scientist