

MIDDAG: Where Does Our News Go? Investigating Information Diffusion via Community-Level Information Pathways

Mingyu Derek Ma¹, Alexander K. Taylor¹, Nuan Wen², Yanchen Liu^{3,4}, Po-Nien Kung¹
Wenna Qin³, Shicheng Wen², Azure Zhou³, Diyi Yang³, Xuezhe Ma², Nanyun Peng¹, Wei Wang¹

¹University of California, Los Angeles

²University of Southern California

³Stanford University

⁴Harvard University

{ma, ataylor2, ponienkung, violetpeng, weiwang}@cs.ucla.edu, {nuanwen, wenshich, xuezhema}@usc.edu
yanchenliu@g.harvard.edu, {wennaqin, amysz, diyiy}@stanford.edu

Abstract

We present MIDDAG, an intuitive, interactive system that visualizes the information propagation paths on social media triggered by COVID-19-related news articles accompanied by comprehensive insights, including user/community susceptibility level, as well as events and popular opinions raised by the crowd while propagating the information. Besides discovering information flow patterns among users, we construct communities among users and develop the propagation forecasting capability, enabling tracing and understanding of how information is disseminated at a higher level. A demo video and more are available at <https://info-pathways.github.io>.

Introduction

The current information propagation ecosystem consisting of traditional news outlets and social media platforms allows for near real-time transmission of information concerning current events. The complexity of the multi-modal data contained in individual information pathway (IP) necessitates the development of novel approaches to analyze how information originally reported by traditional news outlets will spread across social media platforms, as well as the receptivity of the users engaging with this information. Many existing methods perform link prediction on social media networks (Daud et al. 2020; Kumar et al. 2020; Wu et al. 2022; Toprak et al. 2023) and IPs (Taylor et al. 2023; Jin et al. 2023). It is hard to find works that intuitively visualize the existing and forecasted IPs with comprehensive insights on motivating forces of information propagation, such as user/community characteristics and discussion content.

We present a comprehensive visualization system incorporating state-of-the-art information analysis components to provide a clear demonstration of information propagation details and patterns concerning COVID-19-related news within and across the audiences of prominent news organizations from countries highly impacted by COVID-19 across social media platforms. The first dataset consists of all COVID-19-related tweets extracted from the Twitter API. Using tweets containing links to news articles authored by selected news organizations, we construct communities of

users based on their engagement patterns with the selected news organizations, which enables us to construct and visualize the IPs at the community level. We use these pathways to train our novel IP prediction model and present the predicted information propagation patterns for a held-out evaluation set. We also present a novel machine-learning-based approach trained on the collected tweets at the user level to predict and display the susceptibility level of users and communities, which provides further insight into each pathway. In addition to the susceptibility score, we perform event extraction on each tweet in an IP, which shows how and if the core ideas of the original article spread. To diversify the source of our data, we also use the Reddit data available for each news organization and apply existing techniques to extract popular opinions from the available posts in order to provide additional characterization to the IPs once visualized. Finally, we pack the IP analysis capabilities in a system for the user to select, visualize, interact and discover IPs at both user and community levels.

Data and System Design

We use COVID-19-related social media data on Twitter and Reddit to enable a broad IP investigation. We utilize an updated version of the Twitter news dataset presented in prior work (Taylor et al. 2023), including all tweets from May 2020 to April 2021 containing COVID-19 keywords. Due to the large dataset size, we focus on the subset containing tweets from May 15 to May 30, 2020, which contained the largest number of tweets among all periods, including 640 million tweets from 5.3 million users. We consider news articles from the most prominent organizations (according to the Digital News Report from the Reuters Institute) for 15 countries with the most COVID-19 cases as the start of pathways. We retrieve all user-level IPs, including source tweets mentioning a news URL and subsequent retweets and replies. Besides Twitter, we explored discussions on Reddit with hyperlinked COVID-19 news articles and within the same period as the Twitter subset. In total, we collected 5,410 posts on 4,578 unique articles across 649 subreddit communities.

Our system demonstrates IPs and their properties. The user first types keywords to search and picks news articles that serve as the starting points of IPs. The user-level visual-

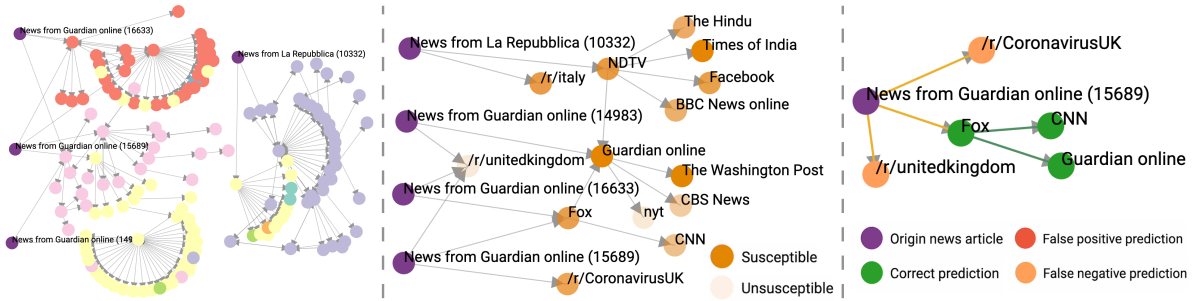


Figure 1: Visualizations of the user-level information pathways (left), community-level information pathways with estimated susceptibility level for each community (center), and community-level information propagation forecasting results (right).

ization includes reply/repost propagation among users, susceptibility score and community assignment. Community-level visualization consists of a directed graph indicating IPs among communities, community aggregated susceptibility levels and key opinions. The event panel shows the event trigger, type and associated arguments in users’ discussions.

Components

We first construct communities among social media users to aggregate the user-level IP to the community granularity. We develop ML models to forecast the information flow, and estimate the susceptibility levels. We further demonstrate events and leading opinions in users’ discussions.

Community Construction We assign users to communities centered around specific news organizations. The community aggregation method follows prior work that measures the influence of nodes in social networks (Romero et al. 2010). It measures the influence that a given node exerts on its neighborhood and the likelihood of how passive or receptive to propagation from its neighborhood a given node is. For each user that interacts with the given community, we retrieve the number of URLs it has posted and compile its one-hop neighborhood based on the source tweets from the community in question it has both posted and interacted with; these individual user graphs are then united to construct a directed community graph. We then assign weights to each user where Q_i is the number of URLs that i mentioned and S_{ij} is the number of URLs mentioned by i and retweeted by j as illustrated in the following equations. Finally, we perform the Influence-Passivity algorithm until the Influence (I_i) and Passivity (P_i) scores have converged.

$$w_e = \frac{S_{ij}}{Q_i}, u_{ij} = \frac{w_{i,j}}{\sum_{k:(k,j) \in E} w_{k,j}}, v_{ji} = \frac{1 - w_{ji}}{\sum_{k:(j,k) \in E} (1 - w_{jk})}; I_i \leftarrow \sum_{j:(i,j) \in E} u_{ij} P_j, P_i \leftarrow \sum_{j:(j,i) \in E} v_{ji} I_j$$

Pathway Prediction We use an improved version of the state-of-the-art IP prediction model (Taylor et al. 2023). We use the community assignments described above to map each user-level information pathway to the community level. The community-level pathways of time periods other than the 15-day selected duration are separated into distinct time windows for training and evaluation, and a graph neural network model is trained to perform link prediction. The evaluation results show that the IP prediction model yields

86.83% AUC for the link prediction task. We apply the trained model to the unseen IP graphs to conduct autoregressive prediction of the full information propagation traces.

Susceptibility Users’ susceptibility reflects their reactions to a piece of misinformation. Since collecting an unobservable belief is hard, we develop an ML model to estimate susceptibility by analyzing its influence on users’ repost behavior (P_{repost}) (Liu et al. 2023). When a user (u) perceives a piece of content (c), a more susceptible user is more likely to repost c . The reposting probability is calculated by

$$P_{repost} = Sigmoid(E(u) * E(c) * Sus(E(u), E(c)))$$

where u is represented by user history posts, E is the text embedding model based on RoBERTa-large (Reimers and Gurevych 2019), and the susceptibility score is produced by the Sus - an MLP with the user and content embeddings as input. We perform contrastive learning to tune the model to distinguish reposted (u, t) pairs from non-reposted ones. We train the ML model with misinformation tweets in the ANTi-Vax (Hayawi et al. 2022) and CoAID (Cui and Lee 2020) datasets and we retrieve corresponding user profiles through the Twitter API. The model produces 86.28% F1 score for retweet behavior prediction, indirectly indicating its reliable performance for susceptibility modeling. To obtain an aggregated susceptibility score for a community, we calculate the mean of individual susceptibility scores.

Event Extraction and Community Opinion Event extraction (EE) aims to identify triggers and arguments for events mentioned in the text, which enables us to understand users’ discussion at a scale (Ma et al. 2021, 2023, 2024). We first define a COVID-19 related event ontology including 9 event types (*i.e.* end organization, social distancing, lock down, quarantine, vaccinate, die, fine, transport and extradite) and their arguments. Since there are no existing EE annotations on COVID-related events, we perform instruction tuning to enable a T5-large model (Raffel et al. 2023) to generalize to newly defined event types following UIE (Lu et al. 2022) by pre-training it with text-structure pairs and further fine-tuning it on 13 datasets of entity/relation/event/sentiment extraction tasks encoded with a unified language. To provide discussion details among users in a community when it propagates certain information, we identify the most liked post among downstream posts in an information flow as the representative opinion of the community.

Acknowledgments

Many thanks to the anonymous reviewers for their feedback. This effort was sponsored by the Defense Advanced Research Project Agency (DARPA) grant HR00112290103.

References

- Cui, L.; and Lee, D. 2020. CoAID: COVID-19 Healthcare Misinformation Dataset. arXiv:2006.00885.
- Daud, N. N.; Ab Hamid, S. H.; Saadon, M.; Sahran, F.; and Anuar, N. B. 2020. Applications of link prediction in social networks: A review. *Journal of Network and Computer Applications*, 166: 102716.
- Hayawi, K.; Shahriar, S.; Serhani, M.; Taleb, I.; and Mathew, S. 2022. ANTi-Vax: A Novel Twitter Dataset for COVID-19 Vaccine Misinformation Detection. *Public Health*, 203: 23–30.
- Jin, Y.; Lee, Y.-C.; Sharma, K.; Ye, M.; Sikka, K.; Divakaran, A.; and Kumar, S. 2023. Predicting Information Pathways Across Online Communities. *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Kumar, A.; Singh, S. S.; Singh, K.; and Biswas, B. 2020. Link prediction techniques, applications, and performance: A survey. *Physica A-statistical Mechanics and Its Applications*, 553: 124289.
- Liu, Y.; Ma, M. D.; Qin, W.; Zhou, A.; Chen, J.; Shi, W.; Wang, W.; and Yang, D. 2023. From Scroll to Misbelief: Modeling the Unobservable Susceptibility to Misinformation on Social Media. arxiv:2311.09630.
- Lu, Y.; Liu, Q.; Dai, D.; Xiao, X.; Lin, H.; Han, X.; Sun, L.; and Wu, H. 2022. Unified Structure Generation for Universal Information Extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5755–5772. Dublin, Ireland: Association for Computational Linguistics.
- Ma, M. D.; Sun, J.; Yang, M.; Huang, K.-H.; Wen, N.; Singh, S.; Han, R.; and Peng, N. 2021. EventPlus: A Temporal Event Understanding Pipeline. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, 56–65. Online: Association for Computational Linguistics.
- Ma, M. D.; Taylor, A.; Wang, W.; and Peng, N. 2023. DICE: Data-Efficient Clinical Event Extraction with Generative Models. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15898–15917. Toronto, Canada: Association for Computational Linguistics.
- Ma, M. D.; Wang, X.; Kung, P.-N.; Brantingham, P. J.; Peng, N.; and Wang, W. 2024. STAR: Improving Low-Resource Information Extraction by Structure-to-Text Data Generation with Large Language Models. *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2023. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. arxiv:1910.10683.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Romero, D. M.; Galuba, W.; Asur, S.; and Huberman, B. A. 2010. Influence and passivity in social media. *Proceedings of the 20th international conference companion on World wide web*.
- Taylor, A. K.; Wen, N.; Kung, P.-N.; Chen, J.; Peng, V.; and Wang, W. 2023. Where Does Your News Come From? Predicting Information Pathways in Social Media. *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Toprak, M.; Boldrini, C.; Passarella, A.; and Conti, M. 2023. Harnessing the Power of Ego Network Layers for Link Prediction in Online Social Networks. *IEEE Transactions on Computational Social Systems*, 10(1): 48–60.
- Wu, H.; Song, C.; Ge, Y.; and Ge, T. 2022. Link Prediction on Complex Networks: An Experimental Survey. *Data Science and Engineering*, 7(3): 253–278.