

RESEARCH ARTICLE

Constructing a Dataset to Support Agent-Based Modeling of Online Interactions: Users, Topics, and Interaction Networks

ABDUL SITTA¹, MIHA ČEŠNOVAR², ALENKA GUČEK¹, AND MARKO GROBELNIK¹

¹E3 Department, Jožef Stefan Institute, 1000 Ljubljana, Slovenia

²Faculty of Mathematics and Physics, University of Ljubljana, 1000 Ljubljana, Slovenia

Corresponding author: Abdul Sittar (abdul.sittar@ijs.si)

This work was supported by TWin of Online Social Networks (TWON) (project number 101095095) as a research project is fully funded by the European Union under the Horizon Europe framework (HORIZON-CL2-2022-DEMOCRACY-01-07).

ABSTRACT Agent-based modeling (ABM) provides a powerful framework for exploring how individual behaviors and interactions give rise to collective social dynamics. However, most ABMs rely on handcrafted or parameterized agent rules that are not empirically grounded, thereby limiting their realism and validation against observed data. To address this gap, we constructed a large-scale, empirically grounded dataset from Reddit to support the development and evaluation of agent-based social simulations. The dataset includes 33 technology-focused, 14 climate-focused, and 7 COVID-related aggregated agents, encompassing around one million posts and comments. Using publicly available posts and comments, we define agent categories based on content and interaction patterns, derive inter-agent relationships from temporal commenting behaviors, and build a directed, weighted network that reflects empirically observed user connections. The resulting dataset enables researchers to calibrate and benchmark agent behavior, network structure, and information diffusion processes against real social dynamics. Our quantitative analysis reveals clear topic-dependent differences in how users interact. Climate discussions show dense, highly connected networks with sustained engagement, COVID-related interactions are sparse and mostly one-directional, and technology discussions are organized around a small number of central hubs. Manual qualitative analysis further shows that agent interactions follow realistic patterns of timing, similarity between users, and sentiment change.

INDEX TERMS Agent-based modeling, online social interactions, information diffusion, network structure analysis, homophily, Reddit discussion networks.

I. INTRODUCTION

Agent-based modeling (ABM) has emerged as a key approach in computational social science for capturing how individual users, each with distinct behavioral rules and network positions, interact to produce collective dynamics. Prior works ([1], [2], [3]) highlighted that ABMs can model heterogeneous agents, complex network structures, emergent phenomena, and support calibration with digital data. More recent studies, such as [4], demonstrate large-scale simulations of information spread across networks, while [5]

The associate editor coordinating the review of this manuscript and approving it for publication was Xiwang Dong.

shows how generative agents based on large language models (LLMs) enhance the realism and adaptability of these models. Despite these advances, a persistent challenge remains: defining realistic agent behaviors based on rich, real-world interaction data and validating simulations against empirical baseline networks.

One of the main challenges in agent-based simulation of online social behavior is the lack of realistic micro-level behavioral baselines. Many ABMs operate at an aggregate level; they rely on parameterized rules for agent posting, commenting, and following, but lack rich empirical interaction data to calibrate these behaviors ([6]). Calibration of ABMs is hindered by large parameter spaces, randomness

in agent behavior, and limited observational data about individual users ([7]). Further, validation remains difficult: without detailed empirical benchmarks of how users follow, comment, and diffuse content, it is hard to evaluate how well a simulation replicates reality ([8]). Table 1 summarizes the key differences between traditional agent-based models and recent LLM-driven ABMs with respect to agent design, empirical grounding, validation strategies, and data requirements have been identified in the literature.

These limitations become even more pronounced as ABMs increasingly incorporate LLM-driven agents, which simulate human-like dialogue, decision-making, and reasoning. While LLM-based agents can exhibit nuanced linguistic and contextual behavior, their outputs can diverge from real-world social dynamics if not empirically grounded ([12], [15], [19]). To address this gap, we construct a large-scale, publicly available dataset from Reddit consisting of posts and comments, which provides a behavioral foundation for grounding and validating both traditional and LLM-based agent simulations. We categorize similar users based on content using topic and interaction features, and represent each such collection of users as a single agent. Instead of relying on explicit follow or friendship data (Reddit does not expose) we estimate directed social relations by analyzing temporal commenting patterns between users. Specifically, when users consistently engage with one another's posts over time, we interpret these interactions as indicative of follower-following-like relations. This approach enables us to construct an empirically grounded network of inter-agent connections derived directly from observed social behavior. The resulting dataset can provide a foundation for calibrating, fine-tuning, and benchmarking both rule-based and LLM-based agent behaviors against real social interaction dynamics.

Social media platforms host domain-specific communities in which users exhibit distinct interaction patterns, linguistic styles, emotional expressions, and information diffusion dynamics. Prior studies demonstrate that behavioral and network characteristics vary substantially across topical domains such as politics, health, climate, and technology, influencing how information spreads and how users engage with content [20], [21]. For example, health and crisis related discussions often display heightened emotional intensity and rapid information cascades, whereas technology or science oriented communities tend to exhibit more sustained, expertise driven engagement [22], [23]. Capturing these domain specific differences is critical for agent-based modeling, as agent behaviors calibrated on a single topic may fail to generalize across domains. Multi domain social media datasets therefore enable the construction of heterogeneous agent populations whose behaviors, interaction frequencies, and network structures reflect real variations in online discourse. Such diversity is essential for developing robust simulations that can model domain dependent phenomena such as misinformation spread, polarization, or collective

attention shifts, and for validating agent based models against empirical baselines drawn from multiple social contexts [24], [25].

We select Reddit as the social ecosystem for this work for several reasons. Hosting hundreds of millions of users across thousands of thematic communities (subreddits), Reddit provides publicly accessible data for both posts and comments, making it one of the richest open social media platforms for large-scale computational social science. For instance, the Pushshift Reddit dataset contains tens of millions of submissions and comments spanning many years, enabling over-time, cross-community, and network-based analyses ([26]). Moreover, Reddit covers a wide range of topics (from politics and science to entertainment and niche hobbies) allowing us to categorize users by content and define meaningful agent types. Its public nature and prior usage in content, network, and demographic analyses ([27], [28]) further support its suitability for mapping user categories and deriving follow/comment relationships. While prior work ([8], [29]) highlights technical and coverage limitations in Reddit data, we address these through careful preprocessing and topic-specific filtering. Overall, Reddit offers a scalable, heterogeneous, and openly accessible environment in which we can ground our agent-based simulation approach particularly for developing and validating generative, LLM-powered agents that model realistic social behaviors.

A. WHY EMPIRICAL DATASETS ARE NEEDED FOR ABMS

The emergence of large language model (LLM)-based agent simulations has underscored a pressing need for empirically grounded datasets that capture real patterns of social interaction, attention, and communication. While recent frameworks (such as EconAgent [30], Generative Agents [10], and AgentSociety [15]) demonstrate that LLM-driven agents can engage in rich, naturalistic dialogues, they largely rely on synthetic or hand-crafted role descriptions and interaction prompts. Without grounding these simulated agents in real-world behavioral distributions, interaction frequencies, and network topologies, such simulations risk producing socially plausible but empirically invalidated dynamics.

In contrast, large-scale social datasets provide the detailed real-world evidence needed to calibrate and validate LLM-based agent societies. Reddit is a particularly valuable source: it contains millions of temporally ordered, user-generated interactions organized around explicit topical communities. These characteristics enable not only the construction of empirically derived social graphs but also the extraction of behavioral baselines such as posting frequency, conversational persistence, and topical engagement intensity that can directly inform agent-level decision parameters in ABMs.

By providing a dataset where users are categorized into topic-aligned agent types and linked via empirically inferred following relationships, we enable a new class of

TABLE 1. Comparison of empirical data requirements and grounding approaches in traditional and LLM-driven agent-based models (ABMs).

Aspect	Traditional ABMs	LLM-driven ABMs
Agent design	Agents are defined by pre-specified behavioral rules and parameters, often handcrafted or calibrated from small-scale empirical studies [2], [9].	Agents are instantiated as Large Language Models (LLMs) with adaptive, generative capabilities that emulate human cognition, dialogue, and decision-making [5], [10].
Data source for calibration	Relies on surveys, small experimental datasets, or theoretical assumptions; difficult to validate at scale [11].	Can be grounded in large-scale text corpora (e.g., Reddit, Twitter, news), offering more realistic social context and variability [12], [13].
Empirical validation	Validation typically performed via statistical comparison of emergent macro-patterns with real-world data [8], [14].	Validation can be performed through behavioral alignment with empirical conversational or interactional datasets [15], [16].
Realism of behavior	Limited by manually defined rules; difficult to capture nuanced linguistic or emotional behavior [17].	Enhanced realism through LLM reasoning and dialogue, provided empirical grounding [5], [10].
Data requirements	Moderate; structured quantitative inputs [18].	High; large-scale unstructured text data [12].
Challenges without empirical data	Poor external validity [8].	Risk of hallucinated or unrealistic behaviors [13].
Empirical grounding benefits	Improved calibration and credibility [9].	Reproducible, data-driven synthetic societies [15].
Example data domains	Mobility, census, epidemiology [14].	Social media, conversation logs [10].

data-aligned simulations. Such datasets allow modelers to (1) fine-tune LLM agents toward realistic social behaviors observed in actual online ecosystems, (2) benchmark the emergent dynamics of simulated societies against observed network metrics, and (3) explore counterfactual or policy experiments in environments constrained by real human behavioral patterns rather than synthetic assumptions. This empirical grounding thus forms the essential bridge between generative simulation and measurable social reality, allowing agent-based models to move from narrative realism toward scientific reproducibility. In this work, an agent does not correspond to a single user but to an aggregated group of users with similar behavioral and content characteristics. This aggregation (macro-agent representation) is a modeling choice that improves simulation efficiency, reduces noise, and enhances interpretability.

II. CONTRIBUTIONS

Our key contributions are as follows:

- We propose an empirically grounded pipeline for building agent-based simulation inputs from large-scale social media data. The pipeline categorizes agents based on their observed behavior. It infers interactions from user activity patterns. It then constructs interaction networks suitable for simulation.
- We release a large-scale dataset derived from Reddit across three topical domains: technology, climate, and COVID-19. The dataset maps users to content-based agent types. It also includes a directed, weighted interaction network inferred from temporal commenting behavior.

- We provide benchmark analyses of user activity and network structure. We also study triadic closure and how it evolves over time. These analyses can be used to calibrate and validate agent-based simulations.

III. RELATED WORK

A. AGENT-BASED MODELING AND SOCIAL SIMULATION

Agent-based modeling (ABM) has become a key approach in computational social science, allowing researchers to study how individual behaviors lead to large-scale social patterns ([2], [31]). Early ABMs, such as the Schelling segregation model and Axelrod's cultural dissemination model, showed how simple local rules can produce complex outcomes. In online settings, ABMs have been used to study the spread of emotions, influence, and information ([4], [32], [33], [34]). While traditional ABMs rely on hand-crafted rules, recent approaches use large language models (LLMs) or data-driven techniques to create more realistic agents that can generate behavior informed by real data ([5], [15]). These advances help bridge the gap between synthetic simulations and real-world social dynamics. Traditional ABMs face persistent limitations, including simplistic decision logic, lack of empirical calibration, and limited scalability or validation ([17], [35]). Manually designed agents often ignore evolving goals, cognitive context, and heterogeneity factors crucial for modeling realistic online interactions. Recent work has sought to address these issues through data-driven and learning-based agents, including LLM-powered agents ([5]), moving toward empirically validated simulations. Our work contributes to this direction by constructing a Reddit-based dataset where users are categorized into agent types derived from content, and inter-agent relationships

are inferred from comments and follows. This provides a data-aligned foundation for simulating and evaluating social dynamics, bridging the gap between manually designed and data-grounded agent-based models.

B. DATA-DRIVEN SIMULATION AND FINE-TUNING OF AGENTS

Recent research has leveraged large language model based agents and empirical data to increase the realism of agent-based simulations. Reference [12] survey the space of LLM-driven social simulation and categorize simulation types from individual to societal scale. The work by [5] introduces LLM-empowered agent frameworks for social network simulation, while [15] build a large-scale system with 10k+ agents interacting in a rich societal environment. Reference [13] emphasis the necessity of grounding agent profiles in real-world data rather than synthetic ones, and [16] use LLM agents in simulated social media to test influence mechanisms. These works collectively point to the potential of combining real interaction data with agent simulation. However, there remains a gap: existing simulations often lack fine-grained mapping of user content interaction data (posts, comments, follows) into agent categories and calibrated inter agent relations. Our work addresses this by collecting large-scale user content/comment data from Reddit, converting users into content-based agent categories, estimating inter-agent relations from comment behavior, and releasing a dataset to enable fine-tuning of agents and evaluation of simulation realism.

C. REDDIT AND ONLINE NETWORK ANALYSIS

Recent research has explored Reddit as a platform for network analysis, community structure, and information diffusion. Reference ([36]) studied the tree structures of Reddit comment threads to understand how global and local features shape discussion depth, width, and size. Reference [37] investigated community identity and its effect on user engagement and retention across multiple subreddits. Reference [38] examined inter-community interactions, mobilizations, and conflicts, highlighting the complex dynamics of Reddit's multi-community environment. Other works, such as ([39]) and ([40]), have constructed subreddit-level networks to study cross-community relations and information flows. Additionally, studies have examined temporal dynamics within user sessions, showing phenomena like performance deterioration over active sessions ([41]). These prior studies provide a strong foundation for network, content, and temporal analyses, yet few integrate content-based user categorization, follower/following relationships, comment interactions, and agent-based simulation, which is the focus of our work. Beyond Reddit, systems such as MID-DAG focus on social media-driven information diffusion, modeling how COVID-19 news propagates through user engagement, event tracking, and diffusion forecasting [42], [43], [44], [45].

D. BRIDGING REAL AND SYNTHETIC SOCIAL DATA

A key challenge in social simulation is ensuring that synthetic or agent-based models meaningfully reflect real-world network and behavioral patterns. Several prior works have addressed this: for example, [46] compared synthetic contact networks with empirical networks, demonstrating the divergence in structural metrics. Similarly, [47] derived agent behavior from an online community and validated an ABM by reproducing the observed network. On the network-generation side, heuristic methods such as personality or compatibility based synthetic networks have attempted to close this realism gap [18]. More broadly, methodological reviews emphasize that combining Social Network Analysis (SNA) and Agent-based modeling (ABM) offers a powerful route to bridge empirical data and simulation frameworks [48]. On the cutting edge, the recent work by [5] integrates generative language-model agents with network simulation and validates against real social media behavior. Our work builds on this trajectory by collecting a large-scale empirical dataset of users, posts, comments from Reddit, converting users into agent categories, estimating inter-agent relations via comment behavior, and releasing a dataset and simulation-ready framework. Unlike many prior studies (which either simulate networks without detailed content-based agent categorization or adapt small/closed communities), our approach combines content, network structure, temporal dynamics, and large-scale publicly accessible data. This dataset can be used to fine-tune agents and benchmark synthetic vs. real discussion dynamics.

IV. METHODOLOGY

Our goal is to construct an empirically grounded foundation for agent-based social simulations using real online interaction data from Reddit. Figure 1 presents the end-to-end methodological pipeline, consisting of five sequential stages: data collection, preprocessing, user categorization and interaction extraction, network construction, and analysis and benchmarking.

We begin by collecting publicly available Reddit data through the Pushshift API or PRAW (Python Reddit API Wrapper). It contains many subreddits, we filtered subreddits named as COVID-19 (2020-2024), technology (2020-2024), and climate changes (2009-2024). Subreddits were selected using a two-stage procedure. First, we applied keyword-based filtering using domain-specific terms (e.g., covid, coronavirus, technology, climate, environment). Second, we manually validated candidate subreddits to ensure topical relevance based on subreddit descriptions and sampled posts.

A subreddit was assigned to a domain if its primary thematic focus aligned with the domain and the majority of its content reflected domain-specific keywords. Subreddits with highly mixed or ambiguous topics were excluded to avoid cross-domain contamination.

In cases where a subreddit covered multiple topics, we assigned it to the domain with the highest content

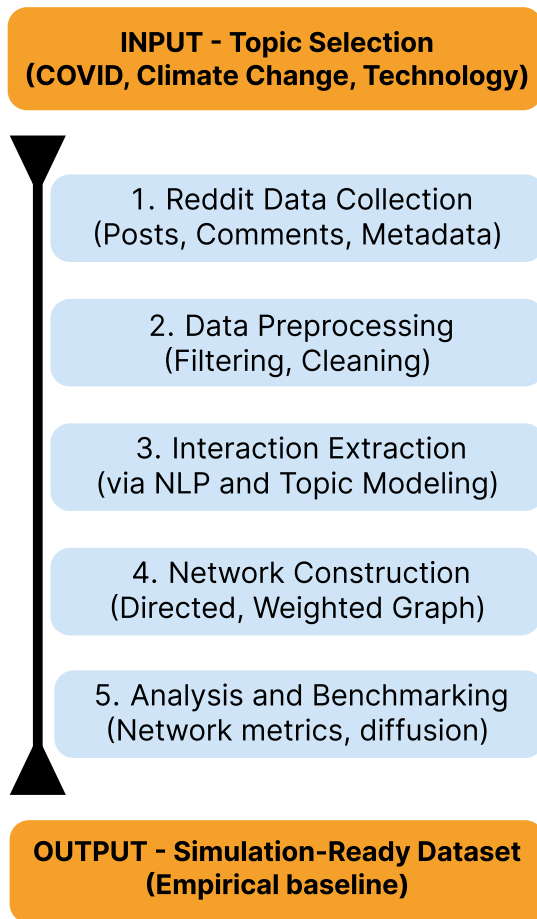


FIGURE 1. Methodological pipeline for constructing an empirically grounded agent-based simulation dataset from Reddit interactions.

overlap based on keyword frequency. The dataset includes posts, comments, timestamps, subreddit identifiers, and anonymized user IDs (see Table 2). This data captures both content (textual data) and structure (who interacts with whom), providing a rich foundation for modeling social behavior (see the Github Link¹). The dataset is publicly available on Zenodo: <https://zenodo.org/records/18082502>²

Raw Reddit data require extensive preprocessing to ensure reliability and reproducibility. Table 2 summarizes the effects of each preprocessing step across the technology, climate, and COVID-19 domains. We:

- Remove deleted or bot accounts and filter out low-quality or spam content.
- Standardize timestamps and metadata.
- Optionally select specific subreddits or topics for focused studies (e.g., politics, gaming, or science).
- Perform language detection and retain only English-language posts, unless multilingual analysis is intended.

¹<https://github.com/abdulsittar/Social-Graph-Inference-Reddit.git>

²<https://zenodo.org/records/18082502>

- These steps ensure a clean, consistent dataset suitable for further semantic and structural analysis.

The substantial reduction in posts, comments, and users across preprocessing steps (see Table 2) reflects a deliberate filtering strategy aimed at retaining only high-quality and interaction-rich data. Specifically, we applied: (1) truncation to the first 10 comments per post to standardize interaction depth and reduce noise from long-tail discussions; (2) removal of bots, deleted content, and low-quality posts; This strict filtering significantly reduces dataset size but ensures that the remaining users exhibit consistent and meaningful interaction patterns. As a result, the final dataset focuses on a small number of highly active users, which are subsequently aggregated into agent types.

To represent heterogeneous agent populations, we categorize users into agent types based on content and behavior. Topic modeling techniques such as BERTopic, and sentence-transformer clustering are applied to users' aggregated posts and comments. Each cluster is then manually assigned an agent category reflecting their primary thematic or behavioral orientation (e.g., informative, opinionated, and supportive). This step grounds agent heterogeneity in empirical data rather than synthetic assumptions.

Next, we infer directed social interactions from comment or reply structures. If user A comments on user B's post or reply, we record a directed edge $A \rightarrow B$. Each edge is weighted by interaction frequency. Temporal ordering is preserved to allow later dynamic modeling. The resulting edges encode the social topology of user interactions observed in the platform.

We aggregate the extracted edges into a directed, weighted graph where nodes represent agents and edges represent empirical interactions. The graph captures realistic degree distributions, clustering, and community structures observed in online environments. Structural features such as reciprocity, triadic closure, and modularity can be analyzed to benchmark the fidelity of the empirical social fabric.

Finally, we compute network and behavioral metrics (e.g., degree centrality, assortativity, clustering coefficient) and simulate diffusion dynamics (e.g., information spread, opinion propagation) to compare empirical patterns with those generated by simulated agents. These analyses serve two purposes: 1) They validate that the constructed dataset reflects authentic social dynamics. 2) They provide empirical baselines for calibrating and evaluating future agent-based models.

The outcome of this pipeline is a simulation-ready dataset consisting of: 1) An agent-agent interaction network grounded in real data., 2) Agent category distributions representing behavioral diversity (see Fig. 2), 3) Benchmark metrics for comparing synthetic simulations to observed online phenomena. This methodology bridges the gap between purely synthetic agent-based modeling and empirically validated social simulation, enabling more realistic exploration of online social systems.

TABLE 2. Dataset preprocessing pipeline across technology, climate, and COVID domains.

ID	Domain	Total Posts	Total Comments	Preprocessing Step	Unique users (posts)	Unique users (comments)
0	Technology	2,207,316	30,804,312	Raw data (no filtering)	556,351	2,422,613
1	Technology	284,294	628,841	Bot removal, topic noise filtering, truncated to first 10 comments	93,811	160,468
2	Technology	135,156	152,516	User activity thresholding (min. interactions)	33	33
3	Technology	111,243	151,060	Removal of deleted/removed posts and comments	33	33
4	Technology	111,243	151,060	Stylistic feature extraction (keywords)	33	33
5	Technology	111,243	151,060	Stylistic and emotional feature enrichment	33	33
6	Technology	111,243	151,060	Inferred follower-following relationships	33	33
0	Climate	226,766	1,914,565	Raw data (no filtering)	35,000	152,844
1	Climate	116,710	341,031	Bot removal, topic noise filtering, truncated to first 10 comments	19,854	42,015
2	Climate	101,662	121,139	User activity thresholding (min. interactions)	14	13
3	Climate	95,997	121,139	Removal of deleted/removed posts and comments	14	13
4	Climate	95,997	121,139	Stylistic and emotional feature enrichment	14	13
5	Climate	95,997	121,139	Agent profile construction (keywords and emotion)	14	13
6	Climate	95,997	121,139	Inferred follower-following relationships	14	13
0	COVID	262,912	3,300,050	Raw data (no filtering)	89,893	220,649
1	COVID	98,067	445,194	Bot removal, topic noise filtering, truncated to first 10 comments	41,273	73,835
2	COVID	83,767	179,951	User activity thresholding (min. interactions)	7	7
3	COVID	83,767	156,817	Removal of deleted/removed posts and comments	7	7
4	COVID	83,767	156,817	Stylistic and emotional feature enrichment	7	7
5	COVID	83,767	156,817	Agent profile construction (keywords and emotion)	7	7
6	COVID	83,767	156,817	Inferred follower-following relationships	7	7

Several studies on reciprocity and link evolution further motivate our temporal framing. Research on predicting reciprocity and triadic closure [51], [52] showed that one-way (parasocial) links often evolve into mutual (reciprocal) ties as users exchange repeated interactions. Our model parallels this by distinguishing between “maybe” and “for sure” followings—interpretable as a behavioral progression from parasocial attention to stable reciprocal engagement. Studies such as [53] on parasocial versus reciprocal relationships highlight that user attributes, network degrees, and behavioral similarity shape whether a directed connection becomes mutual; our rule operationalizes these dynamics in the absence of explicit follow data.

More recent research expands follow and reciprocity prediction to dynamic and generative models. For example, dynamic stochastic block models integrating reciprocity [54] capture how communities evolve and reciprocate connections over time. Our approach offers a simple, data-driven alternative that captures similar temporal patterns using discrete interaction windows. Likewise, link-prediction frameworks that combine heuristic and representation-learning features [54] achieve high predictive accuracy but often sacrifice interpretability. In contrast, our method maintains interpretability while revealing consistent interaction-based follow patterns within evolving discussion networks.

Finally, studies on popularity and follow-back behaviors [55], [56] show that reciprocal following and early adopter effects drive visibility and community growth. Our inference of “for sure” ties can be viewed as identifying those stable, mutually attentive connections that correspond to socially reinforced or influence-generating relationships in other platforms.

While prior research largely focuses on predicting follow or reciprocal edges on platforms such as Twitter or

Google+, our model extends these ideas to environments (e.g., Reddit). By formalizing temporal consistency in directed commenting as a proxy for latent following, we bridge behavioral studies of interaction persistence with network-theoretic research on reciprocity and link formation. This provides a unified, interpretable basis for quantifying user attention and community evolution across topic-specific social ecosystems. Additional details, including the list of subreddits, clustering examples, and annotation guidelines, are provided in the Appendix.

V. HOW REALISTIC ARE THE AGGREGATED AGENTS AND THE ESTIMATION OF NETWORK CONNECTIVITY?

A. QUALITATIVE EVALUATION

The qualitative evaluation focuses on interaction timing, cascade structure, homophily, and sentiment evolution, drawing on empirically grounded interaction chains (see Fig. 3 and Table 3). To complement the quantitative structural, network, and behavioral analyses, we conducted a manual qualitative evaluation of selected agent interactions. The objective of this evaluation is to assess whether the generated agent behavior exhibits plausible human-like conversational dynamics, semantic coherence, and social patterns that are difficult to fully capture using automated metrics alone. The evaluation focuses on interaction timing, cascade structure, profile similarity, interaction escalation, homophily, and sentiment evolution, drawing directly from empirically grounded interaction threads.

In total, we manually evaluated 105 interaction chains across three domains: Technology, Climate Change, and COVID-19. To construct these chains, we first performed topic modeling on each generated post and its associated comments. Pairwise comparisons were then conducted between posts and comments using cosine similarity over

TABLE 3. Manual evaluation questions used to assess agent interactions and emergent social behavior.

Q#	Evaluation Question	Aspects Considered	Analytical Perspective
1	Do agents generate meaningful commentary on posts over time?	Post and comment timing	Cascade analysis; information flow and diffusion
2	Are interacting agents similar in terms of their profiles?	Agent profile characteristics	Engagement life cycle
3	Why do two agents interact more frequently after a specific interaction?	Interaction triggers and context	Content interaction patterns
4	Does the discussion extend among agents with similar characteristics?	Similarity between interacting agents	Homophily
5	Do sentiment patterns align with known behavioral dynamics?	Emotion evolution over time	Sentiment analysis
6	Do interactions exhibit multi-turn conversational exchanges?	Presence of back-and-forth replies	Conversational dynamics
7	Does the topic of discussion remain same across all chains?	Similar topic	Topic modeling and topic drift

their topic representations. When the cosine similarity between two items exceeded a threshold of 0.1, we defined this relationship as a semantic connection. Connections were directed from the earlier node to the later node based on posting timestamps, yielding temporally ordered interaction graphs. If a node exhibited multiple outgoing semantic connections, the corresponding interaction chain was duplicated such that each resulting chain contained only a single incoming and a single outgoing connection per node. This normalization step ensured that each chain represented a linear conversational trajectory rather than a branching structure. The resulting chains were ranked by length, and for manual qualitative evaluation we selected the 35 longest chains per domain, resulting in a total of 105 evaluated chains.

1. Information flow and cascade dynamics: Manual analysis of post–comment timelines reveals that agent interactions follow realistic temporal patterns. Initial posts typically trigger a few of replies shortly after publication, followed by declining engagement over time.

2. Engagement life cycle: Manual comparison of agent profiles including dominant keywords, thematic focus, stylistic and emotional features shows that agents who interact repeatedly tend to share overlapping topical interests and communicative styles. Agents with similar profiles engage more frequently and sustain interactions over longer periods.

3. Content interaction patterns: Across multiple evaluated chains, a clear escalation is observed following the first meaningful exchange between agents. Once two agents engage directly (e.g., via an explicit reply), subsequent interactions become more frequent and more targeted.

4. Homophily and social similarity: Manual analysis confirms that agents tend to interact more with others who are similar to them. Similarity appears in topic interest, expression opinions, and communication styles. This pattern is visible in repeated one on one interactions.

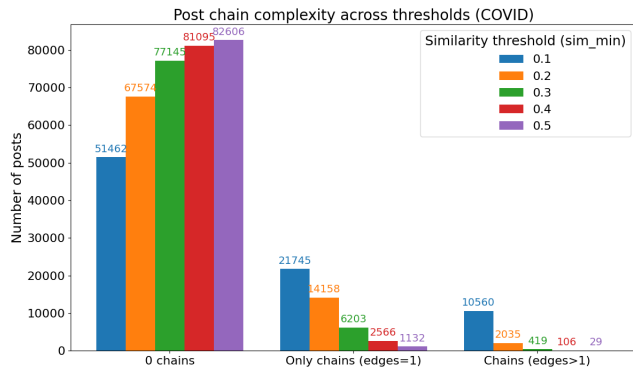
5. Sentiment dynamics: Manual sentiment inspection reveals emotionally coherent trajectories across interaction chains. Neutral or informational exchanges tend to remain emotionally stable, while argumentative or controversial

discussions often show increasing emotional intensity, including frustration or negativity.


6. Conversational dynamics: Extended multi-turn conversations are rare. Most chains consist of only a few turns (typically up to 6–7 comments), often dominated by one or two agents. This limited depth may partly result from agent grouping strategies and similarity-based interaction rules.

7. Topic modeling and topic drift: Topics within interaction chains remain relatively stable. There is little evidence of significant topic drift within short-lived chains.


Evaluation summary: The analysis shows that interaction chains are generally short and limited in scope. Most chains consist of up to 6–7 comments and typically unfold over a few days. Two dominant interaction patterns emerge: one where a user states an opinion and others argue or respond critically, and another where a user asks for similar experiences and others reply in a question–answer or discussion format. About half of the chains involve only a single agent. When multiple agents participate, they are usually very similar in emotional tone and writing style. Often, only one or two aggregated agents dominate the interaction throughout the entire lifecycle of a chain. Multi-turn or extended discussions are rare. When they do occur, they tend to stay focused on the same topic and are usually driven by the same two agents. This limited conversational depth may be influenced by how agents were grouped by topic and similarity. Agent interactions strongly reflect homophily: agents mostly interact with others who are similar to themselves in tone, emotion, and style. Emotional dynamics within chains often show a progression, starting from a neutral tone and gradually intensifying toward stronger emotions such as anger, fear, or joy. These emotional trajectories generally align with both the topic of discussion and the typical emotional profiles of the agents involved, and metadata-based sentiment analysis mostly confirms these patterns. Finally, network growth effects are difficult to identify clearly. However, there is some indication that agents who comment on another agent’s post often end up following that agent, suggesting a weak but observable link between interaction and network connections.



(a) COVID domain: distribution of interaction chains used for qualitative evaluation.



(b) Climate domain: distribution of interaction chains used for qualitative evaluation.



(c) Technology domain: distribution of interaction chains used for qualitative evaluation.

FIGURE 3. Distribution of post-level interaction chain complexity across similarity thresholds (0.1–0.5). For each threshold, posts are grouped by whether they form no interaction chains, only single-edge chains (length = 1), or also longer chains (length > 1). Counts indicate the number of posts in each category.

Overall, the manual qualitative evaluation shows that agent interactions reflect key characteristics of realistic online behavior. These include bursty and cascading interaction patterns, sustained engagement driven by profile similarity, increased interaction after initial contact, homophilic alignment, and believable sentiment changes across multi-turn discussions. Importantly, these observations align well with the quantitative metrics used in our analysis and

help explain them. Together, these results support the conclusion that the generated agent dataset captures not only realistic network structure but also meaningful semantic, emotional, and social behaviors. As a result, the manual evaluation adds an important layer of interpretability and ecological validity, increasing confidence in the dataset’s suitability for agent-based modeling and social simulation studies.

Each interaction chain was manually reviewed by multiple evaluators following the structured criteria outlined in Table 3). In total, 105 interaction chains (35 per domain) were analyzed. Disagreements between evaluators were resolved through discussion to reach a consensus. Each chain typically consisted of multiple posts and comments, allowing assessment of interaction timing, sentiment evolution, and conversational dynamics.

B. QUANTITATIVE EVALUATION

1) IN-DEGREE CENTRALITY: IDENTIFYING INFLUENTIAL NODES

Figure 4a highlights the most prominent users or topics in terms of in-degree, indicating those followed by a large number of others at the 0.01% coverage threshold within the “maybe” category (less certain but still notable connections). The edges represent potential follow relationships estimated from user behaviors such as commenting or engagement patterns. In-degree here measures how many unique users appear to follow or interact with each node, reflecting its popularity or centrality within the network. At the top of the ranking is *DiscordCommunity*, which stands out significantly with the highest in-degree count, suggesting it is the most followed or most connected topic at this detection level. It likely represents a highly active or engaging domain that attracts attention across various user groups.

The network graph (Fig. 4b) displays the top users or topics most followed by others in the COVID domain, based on in-degree at the 0.01% coverage threshold under the “maybe” relationship category. In this context, in-degree reflects the number of inferred connections directed toward each topic, capturing which subjects attract the most user engagement or attention. The most prominent node is *MaskUsage*, which shows the highest in-degree count, suggesting that discussions or user interactions surrounding mask-related behavior were central and widely referenced. Close behind are *DietNutrition* and *SkinSymptoms*, both exhibiting strong connectivity—indicating that users frequently interacted with content linked to health maintenance, nutrition, and symptom tracking during the pandemic period. Further down the ranking are *LongCovid*, *AlternativeTreatments*, and *CovidInfection*, which—while less dominant—still display significant interaction counts. Their presence highlights how extended health concerns and debates about treatment efficacy contributed to sustained user engagement. Collectively, these results show that

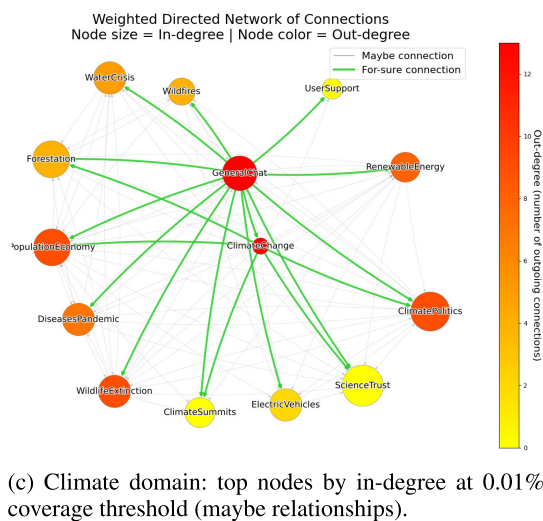
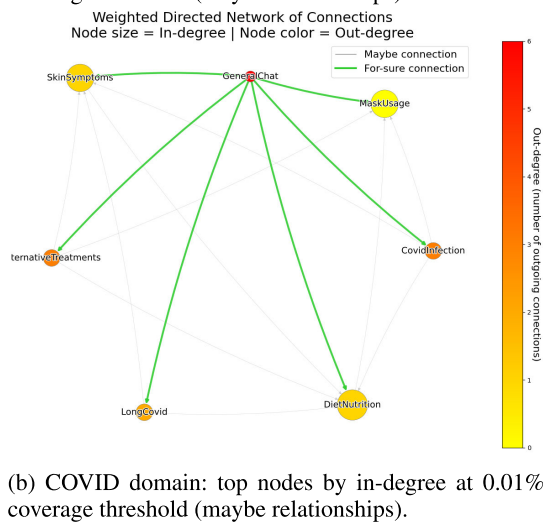
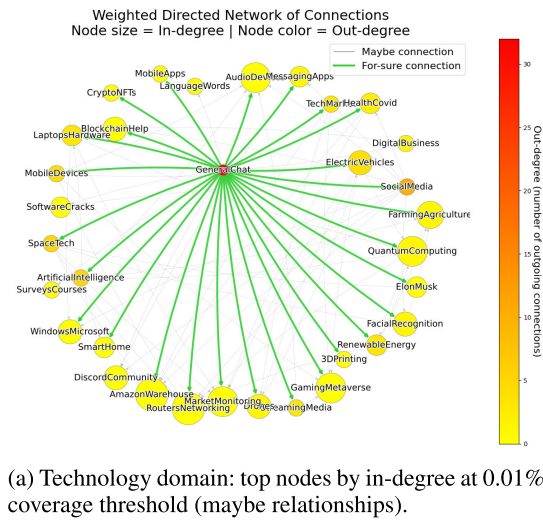


FIGURE 4. Visualization of inferred following relations at a 0.01% similarity threshold across the Technology, COVID, and Climate domains. Edges indicate “for sure” (green) and “maybe” (grey) connections between posts.

health-related subtopics, rather than policy or logistics, dominated attention in the COVID discourse at this inferred

level. The use of the “maybe” network captures a broader behavioral footprint, identifying emerging areas of collective focus where connections may be exploratory or indirect, but nonetheless indicative of influence within the overall information ecosystem.

The network graph (Fig. 4c) presents the top users or themes with the highest inferred in-degree values within the Climate domain at the same 0.01% coverage threshold (“maybe” category). Here, in-degree measures how many other nodes appear to reference or engage with each topic, revealing which areas serve as focal points in climate-related discourse. The leading node, *ScienceTrust*, stands out markedly as the most followed or referenced topic, underscoring the centrality of scientific credibility and public trust within environmental conversations. This is followed by *DiseasesPandemic* and *GeneralChat*, which likely reflect overlap between climate discussions and broader societal or health concerns, as well as general spaces of user interaction. Mid-tier nodes such as *Forestation*, *WildlifeExtinction*, and *ClimatePolitics* indicate active engagement with ecological and political aspects of climate change, suggesting a mix of activism, awareness, and debate-driven interactions. Meanwhile, *Wildfires*, *ElectricVehicles*, *WaterCrisis*, and *UserSupport* occupy the lower range of the top ten but still maintain substantial in-degree values, reflecting persistent but more specialized interest clusters. Overall, the Climate network exhibits a diverse structure of attention, with topics spanning science, policy, and sustainability. The inferred (“maybe”) connections reveal a broad, interconnected community of discussion that, while diffuse, collectively emphasizes trust in science and multi-sectoral environmental awareness.

Across all domains, in-degree centrality reveals distinctive influence structures: the Tech network centers on functional and utility-driven hubs, COVID discussions highlight behavioral and health-oriented engagement, and climate discourse underscores the role of scientific trust and interdisciplinary connectivity.

2) OUT-DEGREE CENTRALITY: IDENTIFYING ACTIVE OR EXPLORATORY NODES

The users or topics with the highest number of outgoing connections are presented in Figure 4a, reflecting strong outward engagement at the 0.01% coverage level for the “maybe” category. In a network sense, out-degree represents how many outgoing links or inferred follows each user or topic initiates—identifying the most active or exploratory nodes in forming connections with others. This visualization captures estimated relationships derived from user interaction data (such as commenting or co-discussion patterns), focusing on potential rather than fully confirmed connections. The chart is dominated overwhelmingly by *GeneralChat*, which exhibits a far higher out-degree count than any other node, suggesting it is the most broadly connected or interactive topic in the dataset. Content or users associated with *GeneralChat* appear to engage with a wide variety of other topics, effectively acting as a central hub or bridge within the network. After

GeneralChat, the next most connected topics—*SocialMedia*, *SpaceTech*, *ElectricVehicles*, *MobileDevices*, and *TechMarket*—have much smaller counts, indicating significantly fewer outgoing links. These nodes still play active roles in forming connections but at a much lower scale compared to *GeneralChat*.


In Figure 4b, the distribution of out-degree identifies the users or topics that most frequently initiate interactions within the COVID-related discussion network under the 0.01% “maybe” threshold. Out-degree in this context represents the number of other users or topics that a node follows, mentions, or interacts with—capturing the extent of its outward engagement. The chart reveals that *GeneralChat* dominates the network, exhibiting the highest out-degree by a significant margin. This suggests that it functions as a highly interactive or bridging node, engaging across numerous COVID-related topics. Following it is *CovidInfection*, which also maintains a strong presence, implying active interaction with various subtopics in the COVID discourse. *AlternativeTreatments* and *LongCovid* occupy mid-tier positions, indicating moderate but targeted engagement—these nodes likely focus on more specific or niche discussions within the larger COVID context.

Figure 4c depicts the most active initiators of connections in the Climate-related network, as measured by out-degree under the 0.01% “maybe” threshold. Here, out-degree reflects how broadly each node reaches out to others, signifying active participation and cross-topic engagement within environmental discussions. In this visualization, *ClimateChange* overwhelmingly dominates, with an out-degree count far surpassing all other nodes. This indicates that *ClimateChange* acts as a major connector and initiator of discussions across the climate discourse, serving as the primary hub linking multiple thematic areas. The next most connected node, *GeneralChat*, also demonstrates significant outward engagement, though still at a much lower scale—suggesting it facilitates broader, cross-domain communication. Below these top two, nodes such as *RenewableEnergy*, *PopulationEconomy*, *WildlifeExtinction*, and *ClimatePolitics* exhibit moderate out-degree values, representing more specialized yet active participants within their respective niches. Toward the bottom of the chart, *DiseasesPandemic*, *WaterCrisis*, *ElectricVehicles*, and *Forestation* maintain low out-degree counts, implying either topic-specific isolation or limited engagement breadth.


Across domains, out-degree centrality highlights different modes of activity: the Tech network is driven by exploratory general-interest nodes, COVID discourse shows limited but health-focused outward engagement, and Climate discussions reveal centralized thematic outreach dominated by broad, unifying concepts like climate change.

3) TRIADIC CLOSURE DYNAMICS: NETWORK COHESION OVER TIME


Next we are interested in the evaluation of triadic closure dynamics across three domains—Technology, COVID, and



(a) Technology domain: triadic closures over time.



(b) COVID domain: triadic closures over time.



(c) Climate domain: triadic closures over time.

FIGURE 5. Comparison of triadic closures over time across all 3 domains for all connections (blue) and only “for sure” connections (orange).

Climate—at a 0.01% coverage threshold. Figure 5 compares the evolution of triadic closures over time, with each panel showing the number of closed triads aggregated in six-month intervals. Triadic closure is a key indicator of network cohesion, representing how fragmented or tightly interconnected a community becomes as indirect relationships (A–B–C) turn into direct ones (A–B–C–A).

Across all three domains, the blue solid line (“All edges = for sure + maybe”) consistently shows an upward trend, indicating an overall increase in network interconnectedness. However, the rate and magnitude of this growth vary by domain.

In the Tech domain (Fig. 5a), triadic closure rises sharply after 2020, growing from near zero to roughly 80 by 2025. This rapid increase suggests an intensifying overlap among users or topics—possibly reflecting how technological discussions and influencers become progressively interlinked over time. The corresponding dashed blue line (“New closures”) fluctuates but remains generally positive, implying steady formation of new triangles. In contrast, the orange (“For sure only”) lines remain flat near zero, showing that confirmed, high-confidence ties contribute minimally to this growth; most of the observed cohesion stems from tentative or inferred relationships.

A comparable but more gradual trend is observed in the COVID domain (Fig. 5b), where the number of closed triads increases slowly, reaching approximately 40 by 2025. This indicates moderate network consolidation—users and topics become more interconnected, but the process is less intense than in the Tech domain. As with Tech, “for sure” closures are nearly absent, implying that while inferred connections are growing, strong mutual links between key nodes remain limited. This could reflect the short-term or event-driven nature of COVID-related discourse, where users cluster temporarily around evolving issues.

In contrast, the Climate domain (Fig. 5c) displays a more gradual yet sustained increase in triadic closure, particularly in the cumulative plots. The total number of closed triads continues to rise steadily across years, indicating long-term structural growth in network cohesion. Interestingly, small but visible steps appear in the “for sure only” lines, suggesting that although fewer in number, high-confidence connections do occasionally form and persist. This pattern may point to a more stable, enduring community structure typical of ongoing policy or scientific discussions, where relationships are built over longer periods rather than through rapid bursts of interaction.

Overall, comparing across domains, Tech exhibits the fastest and largest growth in inferred connectivity, COVID shows moderate but short-lived clustering, and Climate reflects slower yet more stable network formation. The consistent absence of strong triadic closures across all domains underscores that much of the apparent interconnection arises from weaker or uncertain ties, hinting at broad but not deeply cohesive community structures.

4) CROSS-DOMAIN COMPARISON OF INFERRED FOLLOWING DYNAMICS

Figure 6 compares inferred following events across three topical domains—technology, COVID-related health discussions, and climate/environmental topics—using the same consistency-based detection rule defined in Equation 1. In all cases, directed commenting patterns are treated as implicit signals of “following” relationships, where the temporal persistence of such interactions determines their strength.

As shown in Fig. 6a, most strong (for sure) connections emerge during the 2020–2021 period. This




FIGURE 6. Technology domain: detected following events over time at 0.01% threshold. Detected following events over time. Grey dots indicate maybe followings (two or fewer consecutive user interactions), while green dots denote for sure followings (three or more consistent directed interactions) inferred from commenting behavior. The full figures for all three domains have been provided in the appendix.

TABLE 4. Comparative interpretation of interaction networks by topic.

Topic	Key Characteristics
Climate	Dense, highly clustered, moderately reciprocal discussions with sustained interaction and shared attention.
COVID-19	Sparse, low-reciprocity interactions dominated by one-way information exchange.
Technology	Fragmented, hub-driven structure with strong topical segmentation and minimal conversational reciprocity.

concentration suggests a surge of interaction intensity during those years, particularly around discussions involving GeneralChat, SocialMedia, HealthCovid, and technical subtopics such as QuantumComputing and ArtificialIntelligence.

In the COVID dataset (Fig. 6b), early 2020 exhibits the most prominent clusters of confirmed (green) following relations, particularly from GeneralChat toward health-related subtopics such as MaskUsage, LongCovid, and AlternativeTreatments. This temporal clustering corresponds with the outbreak period, when user attention was highly focused and cross-topic engagement peaked.

In contrast, the climate dataset (Fig. 6c) covers a much longer temporal window (2009–2025) and displays several distinct phases of interaction. Between 2010 and 2014, dense clusters of green and grey events occur among ClimateChange, RenewableEnergy, and PopulationEconomy, indicating an early period of sustained cross-topic discussion.

Across domains, a common pattern emerges: large-scale social events or global concerns (e.g., the pandemic, environmental crises, or rapid technological advances) trigger bursts of concentrated engagement that gradually diffuse into smaller, less persistent interaction clusters. The technology and COVID networks show sharper peaks of engagement followed by rapid decline, reflecting event-driven participation. The climate network, by contrast, exhibits longer-term continuity, with engagement sustained across more than a decade but distributed among a broader set of topics.

Together, these results demonstrate that while the inferred following mechanism identifies temporally consistent relationships within each domain, the persistence, density, and thematic breadth of such relationships differ markedly across

social contexts. The method thus provides a comparative lens on how online communities organize around short-term versus enduring societal issues.

VI. ANALYSIS AND BENCHMARKING

This section evaluates the inferred social interaction networks quantitatively, benchmarking their structural properties against established graph-theoretic baselines. The aim is to assess the reliability, robustness, and interpretive validity of the inferred relationships—particularly whether observed patterns such as strong hubs, cohesive clusters, and triadic closures reflect genuine structural organization rather than artifacts of data sparsity or noise. Table 4 summarizes the key structural characteristics of inferred interaction networks across topics, highlighting differences in density, reciprocity, and clustering patterns.

A. NETWORK-LEVEL STRUCTURAL METRICS

We first analyze global structural properties of the inferred interaction networks across topics. Table 5 summarizes core network-level metrics, including density, clustering coefficient, reciprocity, and average path length. Climate discussions produce the densest and most clustered network, indicating sustained and overlapping conversational interactions among users. The high clustering coefficient suggests frequent triadic closures, consistent with conversational cascades and repeated engagement. COVID-related discussions exhibit a sparse and weakly clustered structure, reflecting episodic or broadcast-style interactions. Technology discussions show low density but moderate clustering, indicating selective engagement around central users or technical experts.

Across all topics, the average path length remains low, suggesting efficient information reach despite differing interaction patterns.

TABLE 5. Network-level structural metrics across topics.

Metric	Climate	COVID	Technology
Nodes	14	7	33
Edges	35	7	40
Density	0.192	0.167	0.038
Clustering	0.765	0.295	0.349
Reciprocity	0.286	0.000	0.000
Avg. Path Length	1.67	1.67	1.92

B. DEGREE DISTRIBUTIONS AND ASSORTATIVITY

Degree statistics reveal substantial heterogeneity across topics. Climate interactions exhibit relatively balanced in- and out-degree distributions, suggesting mutual conversational engagement. In contrast, technology discussions are dominated by a small number of high out-degree users interacting with many passive participants, producing hub-centric structures. COVID discussions show minimal degree variance, consistent with limited interaction depth.

The absence of strong assortative mixing across all topics indicates that users do not preferentially interact with others of similar activity levels. Instead, interaction frequency appears driven by topical salience and participation timing rather than structural homophily in degree space.

C. COMMUNITY STRUCTURE AND MODULARITY

We detect communities using modularity maximization and report results in Table 6. Climate discussions form few, weakly separated communities, suggesting a shared conversational space with cross-community engagement. Technology discussions exhibit the highest modularity and the largest number of communities, reflecting topical specialization and segmentation. COVID discussions show limited community formation, primarily due to the small network size rather than strong separation.

These findings suggest that community emergence is topic-dependent, with technical domains encouraging specialization while socio-political topics foster broader interaction.

TABLE 6. Community structure metrics.

Metric	Climate	COVID	Technology
Modularity	0.083	0.122	0.258
Number of Communities	3	2	6
Largest Community Size	7	5	19

D. FILTER BUBBLES AND ECHO CHAMBER TENDENCIES

To quantify exposure isolation, we compute a filter bubble metric based on within-community interaction concentration. Across all topics, this metric remains at or near zero, indicating no strong evidence of echo chambers within inferred interaction networks.

This suggests that although communities exist particularly in technology discussions, they are not isolated. Instead, users maintain interaction pathways across community boundaries, likely facilitated by shared discussion threads and temporal overlap.

E. ROBUSTNESS AND SENSITIVITY ANALYSIS

We assess robustness by varying interaction thresholds and temporal persistence requirements used to infer latent followings. Network-level patterns, including clustering trends and degree heterogeneity, remain stable across reasonable parameter ranges. While absolute edge counts fluctuate, relative topic-level differences persist, indicating that observed structural distinctions are not artifacts of threshold selection.

This stability supports the validity of interaction persistence as a signal for latent social ties.

F. TEMPORAL BENCHMARKING AND STRUCTURAL EVOLUTION

Finally, we examine how inferred network structure evolves over time. Despite substantial differences in post volume across topics, inferred networks stabilize quickly once

sufficient interaction history accumulates. Climate discussions show gradual densification and increasing clustering, while technology discussions stabilize around hub-dominated configurations. COVID discussions show limited temporal evolution due to sparse interaction repetition.

These findings highlight that temporal persistence of interaction—not raw activity volume—is critical for latent network formation.

G. CROSS-DOMAIN COMPARATIVE INSIGHTS

Comparing inferred interaction networks across domains reveals systematic differences in how conversational attention, engagement persistence, and structural cohesion emerge under varying topical contexts.

Climate discussions consistently exhibit the most cohesive interaction structure, characterized by high clustering, moderate reciprocity, and balanced degree distributions. This suggests sustained multi-turn exchanges and conversational cascades, aligning with issues that invite prolonged debate, norm formation, and collective sense-making. The presence of reciprocal ties further indicates a transition from parasocial attention to stable social engagement.

COVID-related discussions, despite high post volume, generate sparse and weakly connected networks with minimal reciprocity and limited community structure. This pattern reflects episodic engagement driven by external events and information dissemination rather than ongoing dialogue. Interactions are predominantly one-directional, suggesting that users react to content rather than form persistent conversational ties.

Technology discussions occupy an intermediate position, combining low network density with relatively high modularity. A small number of high out-degree users dominate interactions, producing hub-centric structures consistent with expert-driven or influencer-mediated communication. Community segmentation is more pronounced, indicating topical specialization and selective engagement rather than broad conversational mixing.

Across all domains, the absence of strong filter bubble effects suggests that topic-driven participation does not necessarily lead to structural isolation. Instead, interaction persistence and degree heterogeneity jointly shape network evolution in a domain-specific manner.

VII. DISCUSSION

A. TOPIC-SPECIFIC INTERACTION DYNAMICS

Our results demonstrate that interaction-based latent following networks exhibit strong topic-dependent structural differences, despite being inferred from the same underlying behavioral rule. Climate discussions form dense and highly clustered interaction networks, indicating sustained conversational engagement and repeated attention among users. In contrast, COVID-19 discussions appear sparse and largely non-reciprocal, suggesting broadcast-style

or episodic information sharing rather than dialogue. Technology discussions occupy an intermediate regime, characterized by hub-dominated and modular structures that reflect expert-driven or problem-solving interactions. These findings indicate that topics semantics strongly shape how attention and interaction persist over time, even when explicit follow mechanisms are absent.

B. INFORMATION DIFFUSION AND CONVERSATIONAL CASCADES

The low average path lengths observed across all topics indicate that information can spread rapidly through inferred interaction networks. However, the mechanism of diffusion differs by topic. In climate discussions, high clustering and moderate reciprocity suggest cascade-like diffusion driven by conversational reinforcement and repeated exchanges. In technology discussions, diffusion is more centralized with a small number of highly active users initiating interactions with many passive participants. COVID-19 discussions exhibit limited cascade depth, consistent with one-off replies rather than multi-turn conversational chains. This highlights that information flow is not only a function of network connectivity, but also of interaction persistence and reciprocity.

C. HOMOPHILY AND COMMUNITY STRUCTURE

The modularity and community structure further reflect homophilic tendencies that vary by topic. Climate discussions show less modularity and few communities, suggesting that users interact across viewpoints within a shared conversational space. Technology discussions, by contrast, display higher modularity and a larger number of communities, consistent with topical segmentation and specialization. COVID-19 discussions show limited fragmentation, but this is largely driven by the small size of the inferred network rather than strong community structure. These patterns suggest that homophily manifests differently depending on whether a topic encourages debate, information seeking, or broadcasting behavior.

D. RECIPROCITY AND ENGAGEMENT LIFE CYCLES

Reciprocity emerges as a key differentiator across topics. Climate discussions exhibit non-zero reciprocity, indicating back-and-forth exchanges and engagement cycles that persist over time. In contrast, the absence of reciprocity in COVID-19 and technology discussions suggests predominantly parasocial or one-directional attention. This aligns with prior work showing that repeated interactions often precede the formation of stable social ties, while one-off interactions rarely evolve into reciprocal relationships.

E. FILTER BUBBLES AND EXPOSURE DIVERSITY

Across all topics, the filter bubbles metric remains close to zero, indicating no strong evidence of exposure isolation within inferred communities. This suggests that, at least




FIGURE 7. Technology domain: detected following events over time at 0.01% threshold.

within the scope of interaction-based following, users are not confined to tightly segregated attention clusters. Instead,

observed community structures appear to emerge from topical or behavioral specialization rather than ideological




FIGURE 8. Climate domain: detected following events over time at 0.01% threshold.

isolation. This findings highlights that filter bubbles may be less pronounced in conversational settings where interaction

requires direct engagement, as opposed to passive content consumption.




FIGURE 9. COVID domain: detected following events over time at 0.01% threshold.



FIGURE 10. Qualitative evaluation (example 1): Cascade analysis, engagement life cycle, homophily, sentiment analysis, conversational dynamics, topic drifting.

F. IMPLICATIONS FOR MODELING SOCIAL ATTENTION

Taken together, these results demonstrate that simple, interpretable rule based on temporal interaction consistency can recover meaningful latent social structures across diverse topics. The inferred networks capture conversational depth, engagement asymmetry, and community organization without relying on explicit social links. This makes the approach particularly suitable for platforms such as Reddit, where following relationships are implicit and topic-dependent.

VIII. VALIDATION AND LIMITATIONS

We validate our inference framework through multiple complementary strategies. First, robustness analyses show that network-level trends remain stable under variations in temporal windowing and interaction thresholds. Second, manual evaluation by human annotators confirms that repeated directed interactions often correspond to meaningful conversational attention and engagement continuity.

Despite these strengths, several limitations remain. The inference of following relationships relies solely on observable reply behavior and does not account for unobserved attention mechanisms such as passive reading or external platform interactions. Temporal aggregation may also obscure short-lived but meaningful interactions, particularly in fast-moving domains such as COVID-related discussions. Furthermore, the relatively small inferred networks—especially in narrower domains—limit the detection of higher-order structural phenomena such as rich-club effects or long-range diffusion pathways.

Finally, while the framework is designed to be platform-agnostic, domain-specific norms and moderation practices may influence interaction visibility and persistence. As a result, inferred following ties should be interpreted as proxies for sustained attention rather than direct equivalents of explicit social connections.

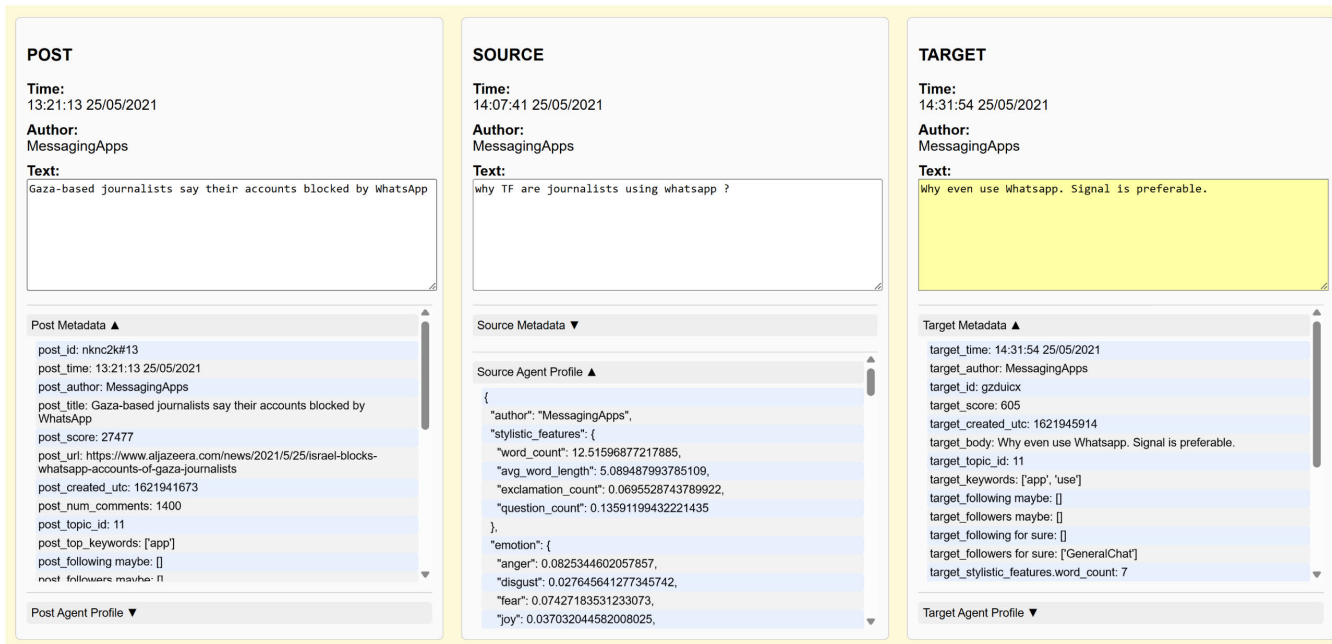


FIGURE 11. Qualitative evaluation (example 2): Cascade analysis, engagement life cycle, homophily, sentiment analysis, conversational dynamics, topic drifting.

Overall, these findings support the validity of temporal interaction persistence as an interpretable signal for latent social ties, while highlighting the need for cautious interpretation and future extensions incorporating content semantics and richer temporal dynamics.

IX. APPLICATIONS AND FUTURE WORK

The inferred interaction networks and benchmarking framework developed in this study provide a versatile foundation for a range of applied and research-oriented use cases. Beyond describing the structural properties of social discussions, these networks serve as testbeds for evaluating social media dynamics, intervention strategies, and computational models of user behavior.

A. APPLICATIONS

The inferred networks enable the simulation and assessment of information interventions, such as moderation policies, content recommendation changes, or targeted awareness campaigns. By modeling how information or influence propagates through empirically grounded structures, researchers can test how modifying connectivity or interaction probabilities might reduce polarization, enhance information diversity, or mitigate misinformation spread. For example, nodes corresponding to bridging topics (e.g., GeneralChat, ScienceTrust) could be used as insertion points for fact-checking or cross-community content to evaluate their diffusion effects under controlled conditions.

The network structure provides a realistic framework on which to create and study simulated social systems. These environments can be used to study emergent

behaviors such as opinion clustering, or viral content diffusion. Because the networks exhibit empirically observed features—small-world structure, and modular clustering—they provide a more realistic base for testing theories of collective behavior, coordination, and conflict formation. Agent-based simulations initialized on these structures can help explore how small changes, such as exposure to new information sources, alter long-term community alignment or fragmentation.

The dataset and derived interaction graphs can further function as a benchmark for evaluating the behavioral realism of social agents or generative dialogue models. By comparing simulated agent interactions against empirical metrics (e.g., reciprocity, clustering, modularity, and degree distributions), researchers can assess whether artificial agents replicate human-like participation patterns or exhibit implausible dynamics. Such benchmarking could guide the development of more sociologically grounded AI systems, ensuring that synthetic communities and conversational agents align with observed social structures.

B. FUTURE WORK

While the current framework provides strong empirical grounding and analytical capability, several directions offer opportunities for enhancement and expansion. Future work will extend the methodology beyond Reddit to platforms such as Twitter/X, Mastodon, and specialized forums. Cross-platform comparison would enable the study of how interaction mechanisms and platform affordances shape network structure, reciprocity, and polarization, offering a more comprehensive view of digital social ecosystems.

Incorporating machine learning models that capture temporal dependencies—for instance, recurrent or dynamic graph neural networks—would allow automated tracking of evolving communities, emerging influencers, and polarization trajectories over time. Modeling cross-community interactions could also reveal how information or users migrate between topics, providing a richer understanding of discourse diffusion and the breakdown (or reinforcement) of filter bubbles.

An important long-term objective is to develop an open simulation environment based on the inferred networks. Such an environment would integrate realistic agent behaviors, communication dynamics, and empirical network structures, serving as a shared research platform for testing interventions, evaluating AI moderation tools, and studying large-scale social phenomena. By releasing these datasets and simulation modules under open standards, future research can systematically benchmark models of collective intelligence, social influence, and misinformation resilience.

Ultimately, integrating these approaches could bridge computational social science and AI research, enabling interpretable, data-driven understanding of online social dynamics. Through continual refinement—adding temporal, semantic, and cross-platform perspectives—the framework can evolve into a unified testbed for studying digital society at scale.

X. CONCLUSION

This work presented the first in combining Reddit-based, agent-aligned dataset designed explicitly for grounding and benchmarking social simulations. By systematically extracting posts, comments, and inferred user-to-user interactions, we constructed a directed, weighted social graph that captures realistic structural and behavioral dynamics across multiple topical domains. Through topic modeling and behavioral clustering, users were categorized into interpretable agent types—allowing the empirical characterization of how different communities interact, follow, and respond to one another.

Our analysis revealed clear empirical regularities: degree distributions followed a heavy-tailed pattern typical of real social networks, while clustering and modularity analyses indicated the presence of well-defined subcommunities. These findings demonstrate that social interactions on Reddit are not randomly distributed but structured around tightly knit, topic-specific clusters that limit cross-domain exposure. Despite these structural constraints, bridging nodes such as GeneralChat, ScienceTrust, and ClimateChange were found to play a disproportionately large role in connecting otherwise fragmented communities—suggesting potential leverage points for intervention or content diversification.

From a modeling perspective, the dataset provides a robust empirical foundation for both traditional and LLM-driven agent-based models. It enables simulation designers to calibrate behavioral rules, validate diffusion outcomes, and benchmark emergent dynamics against real-world interaction

baselines. The results demonstrate that data-grounded simulation frameworks can reproduce observed engagement and connectivity patterns, marking a step toward socially realistic, empirically validated agent societies.

In summary, this study contributes (1) a large-scale, empirically derived dataset linking user content and interaction structure on Reddit, (2) a methodological pipeline for constructing and analyzing social networks suitable for simulation calibration, and (3) empirical insights into online community organization, engagement asymmetries, and filter bubble formation. Together, these elements establish a clear path toward data-grounded, reproducible, and socially valuable agent-based models—bridging the gap between synthetic simulation and real human behavior in online ecosystems.

APPENDIX ACKNOWLEDGMENT

The authors acknowledge the use of artificial intelligence tools in the preparation of this manuscript. Specifically, ChatGPT (OpenAI) was used to assist with drafting and refining portions of the text and with providing explanations related to software implementation. All AI-generated content was reviewed, validated, and approved by the authors, who take full responsibility for the accuracy and integrity of the manuscript.

XI. CONFLICT OF INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article.

REFERENCES

- [1] M. Steinbacher, M. Raddant, F. Karimi, E. Camacho Cuenca, S. Alfarano, G. Iori, and T. Lux, “Advances in the agent-based modeling of economic and social behavior,” *Social Netw. Bus. Econ.*, vol. 1, no. 7, p. 99, Jul. 2021.
- [2] F. Bianchi and F. Squazzoni, “Agent-based models in sociology,” *Wiley Interdiscipl. Rev. Comput. Statist.*, vol. 7, no. 4, pp. 284–306, 2015.
- [3] S. V. Kovalchuk, C. Li, and O. V. Kubryak, “Towards explaining emergent behavior in multi-agent systems micro-parameter space structuring with feature importance in heatbug model,” in *Proc. IEEE/WIC Int. Conf. Web Intell. Intell. Agent Technol. (WI-IAT)*, Dec. 2024, pp. 824–829.
- [4] G. Murić, A. Tregubov, J. Blythe, A. Abeliuk, D. Choudhary, K. Lerman, and E. Ferrara, “Large-scale agent-based simulations of online social networks,” *Auto. Agents Multi-Agent Syst.*, vol. 36, no. 2, p. 38, Oct. 2022.
- [5] A. Ferraro, A. Galli, V. La Gatta, M. Postiglione, G. M. Orlando, D. Russo, G. Riccio, A. Romano, and V. Moscato, “Agent-based modelling meets generative AI in social network simulations,” in *Proc. Int. Conf. Adv. Social Netw. Anal. Mining*. Cham, Switzerland: Springer, 2024, pp. 155–170.
- [6] P. Avegliano and J. S. Sichman, “Equation-based versus agent-based models: Why not embrace both for an efficient parameter calibration?” *J. Artif. Societies Social Simul.*, vol. 26, no. 4, p. 3, 2023.
- [7] J. McCulloch, J. Ge, J. A. Ward, A. Heppenstall, J. G. Polhill, and N. Malleon, “Calibrating agent-based models using uncertainty quantification methods,” *J. Artif. Societies Social Simul.*, vol. 25, no. 2, p. 1, 2022.
- [8] A. Heppenstall, A. Crooks, N. Malleon, E. Manley, J. Ge, and M. Batty, “Future developments in geographical agent-based models: Challenges and opportunities,” *Geographical Anal.*, vol. 53, no. 1, pp. 76–91, Jan. 2021.
- [9] V. Grimm, E. Revilla, U. Berger, F. Jeltsch, W. M. Mooij, S. F. Railsback, H.-H. Thulke, J. Weiner, T. Wiegand, and D. L. DeAngelis, “Pattern-oriented modeling of agent-based complex systems: Lessons from ecology,” *Science*, vol. 310, no. 5750, pp. 987–991, Nov. 2005.

- [10] J. S. Park, J. C. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, "Generative agents: Interactive simulations of human behavior," in *Proc. ACM Symp. User Interface Softw. Technol. (UIST)*, 2023, pp. 1–22.
- [11] T. Lux and M. Marchesi, "Scaling and criticality in a stochastic multi-agent model of a financial market," *Nature*, vol. 397, no. 6719, pp. 498–500, Feb. 1999.
- [12] X. Mou, X. Ding, Q. He, L. Wang, J. Liang, X. Zhang, L. Sun, J. Lin, J. Zhou, and X. Huang, "From individual to society: A survey on social simulation driven by large language model-based agents," 2024, *arXiv:2412.03563*.
- [13] M. Ta, L. Nguyen, and Y. Chen, "SoAgent: A real-world data empowered agent pool to facilitate LLM-driven generative social simulation," in *Proc. VLDB Workshops*, 2025, pp. 45–57.
- [14] R. Hinch, "OpenABM-COVID19—An agent-based model for non-pharmaceutical interventions against COVID-19 including contact tracing," *PLOS Comput. Biol.*, vol. 17, no. 7, Jul. 2021, Art. no. e1009146.
- [15] X. Piao, H. Chen, and M. Li, "AgentSociety: Large-scale simulation of LLM-driven generative agents advances understanding of human behaviors and society," in *Proc. AAAI Conf. Artif. Intell.*, 2025, pp. 1–12.
- [16] B. Jin and W. Guo, "Synthetic social media influence experimentation via an agentic reinforcement learning large language model bot," *J. Artif. Societies Social Simul.*, vol. 28, no. 3, p. 6, 2025.
- [17] R. Conte, N. Gilbert, G. Bonelli, C. Cioffi-Revilla, G. Deffuant, V. Loreto, H. S. Moat, J.-P. Nadal, A. Sanchez, A. Nowak, A. Flache, M. S. Miguel, and D. Helbing, "Manifesto of computational social science," *Eur. Phys. J. Special Topics*, vol. 214, no. 1, pp. 325–346, 2012.
- [18] D. A. O'Neil and M. D. Petty, "Heuristic methods for synthesizing realistic social networks based on personality compatibility," *Appl. Netw. Sci.*, vol. 4, no. 1, p. 19, Dec. 2019.
- [19] F. Rezazadeh, A. Ashtari Gargari, S. Lagen, H. Song, D. Niyato, and L. Liu, "Toward generative 6G simulation: An experimental multi-agent LLM and Ns-3 integration," 2025, *arXiv:2503.13402*.
- [20] V. Lampos, N. Aletas, D. Preotiuc-Pietro, and T. Cohn, "Predicting and characterising user impact on Twitter," in *Proc. 14th Conf. Eur. Chapter Assoc. Comput. Linguistics*, 2014, pp. 405–413.
- [21] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, "The rise of social bots," *Commun. ACM*, vol. 59, no. 7, pp. 96–104, Jun. 2016.
- [22] D. Hayes, "Social media and mental health," *Children Young People Now*, vol. 2018, no. 11, p. 16, 2018.
- [23] W. Wu, H. Lyu, and J. Luo, "Characterizing discourse about COVID-19 vaccines: A Reddit version of the pandemic story," *Health Data Sci.*, vol. 2021, Jan. 2021, Art. no. 9837856.
- [24] D. M. J. Lazer, M. A. Baum, and Y. Benkler, "The science of fake news," *Science*, vol. 359, no. 6380, pp. 1094–1096, 2018.
- [25] C. Sáenz-Royo, C. Gracia-Lázaro, and Y. Moreno, "The role of the organization structure in the diffusion of innovations," *PLoS ONE*, vol. 10, no. 5, May 2015, Art. no. e0126076.
- [26] J. Baumgartner, S. Zannettou, B. Keegan, M. Squire, and J. Blackburn, "The pushshift Reddit dataset," in *Proc. Int. AAAI Conf. Web Social Media*, vol. 14, 2020, pp. 830–839.
- [27] E. Turcan and K. McKeown, "Dreaddit: A Reddit dataset for stress analysis in social media," in *Proc. 10th Int. Workshop Health Text Mining Inf. Anal. (LOUHI)*, 2019, pp. 97–107.
- [28] M. G. M. K. I. Vukojevic and M. B. J. Šnajder, "PANDORA talks: Personality and demographics on Reddit," in *Proc. SocialNLP*, 2021, p. 123.
- [29] D. Gaffney and J. N. Matias, "Caveat emptor, computational social science: Large-scale missing data in a widely-published Reddit corpus," *PLoS ONE*, vol. 13, no. 7, Jul. 2018, Art. no. e0200162.
- [30] X. Li, Y. Xu, Y. Zhang, and E. C. Malthouse, "Large language model-driven multi-agent simulation for news diffusion under different network structures," 2024, *arXiv:2410.13909*.
- [31] N. Gilbert, *Agent-Based Models*. Newbury Park, CA, USA: Sage, 2008.
- [32] M. Šuvakov, D. Garcia, F. Schweitzer, and B. Tadić, "Agent-based simulations of emotion spreading in online social networks," 2012, *arXiv:1205.6278*.
- [33] G. Bernstein and K. O'Brien, "Stochastic agent-based simulations of social networks," in *Proc. 46th Annu. Simul. Symp.*, 2013, pp. 1–8.
- [34] H. Übler and S. Hartmann, "Simulating trends in artificial influence networks," *J. Artif. Societies Social Simul.*, vol. 19, no. 1, p. 2, 2016.
- [35] M. Zhang, A. Verbraeck, R. Meng, and X. Qiu, "Agent-based modeling of large-scale complex social interactions," in *Proc. 3rd ACM SIGSIM Conf. Princ. Adv. Discrete Simul.*, Jun. 2015, pp. 197–198.
- [36] Y. Yu, J. Jiang, and P. S. Dhillon, "Characterizing the structure of online conversations across Reddit," *Proc. ACM Hum.-Comput. Interact.*, vol. 8, no. 1–23, Nov. 2024.
- [37] J. Zhang, W. Hamilton, C. Danescu-Niculescu-Mizil, D. Jurafsky, and J. Leskovec, "Community identity and user engagement in a multi-community landscape," in *Proc. Int. AAAI Conf. Web Social Media*, vol. 11, no. 1, 2017, pp. 377–386.
- [38] S. Kumar, W. L. Hamilton, J. Leskovec, and D. Jurafsky, "Community interaction and conflict on the web," in *Proc. World Wide Web Conf. World Wide Web (WWW)*, 2018, pp. 933–943.
- [39] J. Sawicki, M. Ganzha, M. Paprzycki, and Y. Watanobe, "Reddit CrosspostNet—Studying Reddit communities with large-scale crosspost graph networks," *Algorithms*, vol. 16, no. 9, p. 424, Sep. 2023.
- [40] J. Sawicki and M. Ganzha, "Exploring Reddit community structure: Bridges, gateways and highways," *Electronics*, vol. 13, no. 10, p. 1935, May 2024.
- [41] P. Singer, E. Ferrara, F. Kooti, M. Strohmaier, and K. Lerman, "Evidence of online performance deterioration in user sessions on Reddit," *PLoS ONE*, vol. 11, no. 8, Aug. 2016, Art. no. e0161636.
- [42] M. D. Ma, A. K. Taylor, N. Wen, Y. Liu, P.-N. Kung, W. Qin, S. Wen, A. Zhou, D. Yang, and X. Ma, "MIDDAG: Where does our news go? Investigating information diffusion via community-level information pathways," in *Proc. AAAI Conf. Artif. Intell.*, 2024, vol. 38, no. 21, pp. 23811–23813.
- [43] A. Sittar, D. Mladenčić, A. Guček, and M. Grobelnik, "BAR-analytics: A web-based platform for analyzing information spreading barriers in news: Comparative analysis across multiple barriers and events," *Int. J. Inf. Manage. Data Insights*, vol. 5, no. 2, Dec. 2025, Art. no. 100368.
- [44] A. Sittar, D. Major, C. Mello, D. Mladenčić, and M. Grobelnik, "Political and economic patterns in COVID-19 news: From lockdown to vaccination," *IEEE Access*, vol. 10, pp. 40036–40050, 2022.
- [45] A. Sittar, D. Mladenčić, and M. Grobelnik, "Analysis of information cascading and propagation barriers across distinctive news events," *J. Intell. Inf. Syst.*, vol. 58, no. 1, pp. 119–152, Feb. 2022.
- [46] K. Karra, S. Swarup, and J. Graham, "An empirical assessment of the complexity and realism of synthetic social contact networks," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2018, pp. 3959–3967.
- [47] C. S. Ang and P. Zaphiris, "Simulating social networks of online communities," in *Proc. Int. Conf. Human-Computer Interact.* Cham, Switzerland: Springer, 2009, pp. 198–207.
- [48] M. Will, J. Groeneveld, K. Frank, and B. Müller, "Combining social network analysis and agent-based modelling to explore dynamics of human interaction: A review," *Socio-Environ. Syst. Model.*, vol. 2, p. 16325, Feb. 2020.
- [49] C. J. Hutto, S. Yardi, and E. Gilbert, "A longitudinal study of follow predictors on Twitter," in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, Apr. 2013, pp. 821–830.
- [50] J. Li, J. Tang, Y. Wang, Y. Wan, Y. Chang, and H. Liu, "Understanding and predicting delay in reciprocal relations," in *Proc. World Wide Web Conf. World Wide Web (WWW)*, 2018, pp. 1643–1652.
- [51] J. Cheng, D. M. Romero, B. Meeder, and J. Kleinberg, "Predicting reciprocity in social networks," in *Proc. IEEE 3rd Int. Conf. Privacy, Secur., Risk Trust IEEE 3rd Int. Conf. Social Comput.*, Oct. 2011, pp. 49–56.
- [52] T. Lou, J. Tang, J. Hopcroft, Z. Fang, and X. Ding, "Learning to predict reciprocity and triadic closure in social networks," *ACM Trans. Knowl. Discovery From Data*, vol. 7, no. 2, pp. 1–25, Jul. 2013.
- [53] N. Z. Gong and W. Xu, "Reciprocal versus parasocial relationships in online social networks," *Social Netw. Anal. Mining*, vol. 4, no. 1, p. 184, Dec. 2014.
- [54] H. Safdari, M. Contisciani, and C. De Bacco, "Reciprocity, community detection, and link prediction in dynamic networks," *J. Phys., Complex.*, vol. 3, no. 1, Mar. 2022, Art. no. 015010.
- [55] T. Elmas, M. Randl, and Y. Attia, "# TeamFollowBack: Detection & analysis of follow back accounts on social media," in *Proc. Int. AAAI Conf. Web Social Media*, vol. 18, 2024, pp. 381–393.
- [56] D. Imamori and K. Tajima, "Predicting popularity of Twitter accounts through the discovery of link-propagating early adopters," in *Proc. 25th ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2016, pp. 639–648.



ABDUL SIT TAR received the M.S. degree from the COMSATS Institute of Information Technology, Lahore, Pakistan. He is currently a Postdoctoral Researcher with a specializing in machine learning and natural language processing (NLP), with over a decade of expertise in the field. His work focuses on developing advanced models using frameworks, such as PyTorch, TensorFlow, and Hugging Face. He actively collaborate with academia, contributing to both theoretical and applied research, with a strong record of publications in peer-reviewed international journals. He was a MSCA fellow working on CLEOPATRA ITN.



MIHA ČEŠNOVAR received the bachelor's degree in mathematics and computer science. He is currently pursuing the master's degree in mathematics and computer science with the Faculty of Mathematics and Physics (FMF), University of Ljubljana. He is also a Student Researcher with the Jožef Stefan Institute (JSI), where he assists in research on project TWON. His research interests include the application of mathematical and computational approaches to real-world research problems.



ALENKA GUČEK received the Ph.D. degree in biosciences from the University of Ljubljana, Slovenia. She was a Postdoctoral Researcher in diabetes and molecular cell biology with Uppsala University, Sweden, where she still teaches data visualization. She is currently a professional Research Associate with the Department of Artificial Intelligence, Jožef Stefan Institute (JSI). She is also with the AI Department, JSI, includes codesign of exploratory data visualization tools, observatories, and AI services that meet the needs of users.



MARKO GROBELNIK is currently an Expert Researcher in the field of artificial intelligence (AI). He co-leads the Department for Artificial Intelligence, Jožef Stefan Institute; co-founded the UNESCO International Research Center on AI (IRCAI); and the CEO of Quintelligence.com, specialized in solving complex AI tasks for the commercial world. He collaborates with major European academic institutions and major industries, such as Bloomberg, British Telecom, European Commission, Microsoft Research, and New York Times. He is the co-author of several books, the co-founder of several start-ups and is/was involved into over 50 EU funded research projects in various fields of artificial intelligence. He focused areas of expertise are machine learning, data/text/web mining, network analysis, semantic technologies, deep text understanding, and data visualization.

...