NetSciX2026: International School and Conference on Network Science
February 17^{th}-20^{th}, 2026 - Auckland, New Zealand

Modeling Opinion Dynamics and Social Network Algorithms

Keywords: Opinion dynamics; Adaptive network; Recommendation algorithms; Polarization; Information cascades

Extended Abstract

Social media platforms influence collective opinions by determining what users see and enabling them to reorganize their social ties. We study the coevolution of information flow and network structure using an agent-based model on adaptive, directed networks introduced in [1, 2]. Agents have continuous opinions on a scale of $b \in [-1, 1]$. They can post and repost content from their feeds, and their neighbors receive posts through a recommendation mechanism that filters by topical proximity on the platform. Next, all agents receiving the post update their opinions, and can rewire their connections when they disagree. This model aims to capture key behavioral frictions, such as limited attention and bounded memory, while allowing the platform's delivery policy to be adjusted, which includes "priority" accounts. Specifically, posts from priority users are always delivered to their followers. From the agents' perspective, the model also allows for the simulation of stubborn agents (i.e., agents that do not change their opinions).

To quantify and compare opinion polarization and division in social networks across scenarios, we introduce a homophilic exposure measure. For all agents, we map the pair $(b,\ b_{NN})$, where b_{NN} is the mean opinion among the agent's out-neighborhood [3]. Division appears as mass concentrated in the first and third quadrants of this joint distribution, indicating agreement in sign between the agent and its neighbors. We summarize this structure with a homophilic bimodality coefficient BC_{hom} [4]. To calculate this, we rotate the (b, b_{NN}) cloud by 45° to align its principal axis and compute a bimodality coefficient on the leading dimension (see Figure 1(a)). A higher BC_{hom} indicates stronger polarization into two like-minded groups.

Using this model, we first examine delivery policies that give priority to certain accounts. Specifically, we examine two posting regimes: (i) an agreement-seeking regime, in which agents are more likely to share content similar to their own opinions, and (ii) a reactive regime, in which agents share content with which they disagree. We find that increasing the proportion of priority accounts generally weakens polarization and division (i.e., lower BC_{hom} values). Under reactive posting, polarization breaks down at around a 20% proportion of priority accounts. The network then settles into a more contiguous opinion field with increased cross-cutting exposure. These results suggest that prioritizing certain users can mitigate polarization when reactions such as quote tweets and replies are prevalent.

Next, we examine the dual-use risks. What if some priority account holders are ideologues (i.e., stubborn users with a priority account) with fixed opinions at -1 or $+1$? Introducing even a small percentage of ideologue users reverses the mitigation effect, increasing BC_{hom}. While more priority accounts usually reduce polarization, prioritizing a small percentage of ideologues can substantially raise BC_{hom}, see Figure 1(b).

We also explore the reposting information, such as feed truncation and memory limits on what users can recall for reposting, as proposed in [5]. The reposting mechanism restricts memory capacity and creates new dynamic states. These states include regimes where opinions linger near stable plateaus before changing and regimes where polarization can persist or

reverse, depending on posting rules and the strictness of delivery filters. These experiments clarify how modest changes in exposure breadth and user memory can cause the system to transition between qualitatively distinct outcomes.

Finally, we demonstrate the model's external validity by calibrating it to reproduce the statistical organization of information cascades observed in Twitter data. An information cascade is defined as a set of posts produced by reposting the same piece of content across users. Its size is defined as the number of users who share that content. Using a genetic algorithm, we optimize the dynamics parameters to fit the cascade size. The calibrated model matches the cascade size without relying on individual-level records. This shows that a small set of interpretable mechanisms (i.e., recommendation, rewiring, limited memory, and reposting) suffices to produce realistic cascade statistics while generating the polarization (see Figures 1(c) and (d)). Together, the results support the ability of the opinion model to reproduce social networks to some extent.

This study is simulation-based. The empirical component fits aggregate cascade statistics, and no personal data are collected or inferred, nor individual targeting is performed. We analyse dual-use risks (e.g., the potential for prioritization to amplify polarized minorities) and position the framework for auditing and stress-testing policy changes rather than optimizing manipulative interventions. The methods are not designed for policing or military targeting.

References

[1] Henrique F de Arruda et al. "Modelling how social network algorithms can influence opinion polarization". In: *Information Sciences* 588 (2022), pp. 265–278.

[2] Henrique F de Arruda, Kleber Andrade Oliveira, and Yamir Moreno. "Echo chamber formation sharpened by priority users". In: *iScience* 27.11 (2024).

[3] Wesley Cota et al. "Quantifying echo chamber effects in information spreading over political communication networks". In: *EPJ Data Science* 8.1 (2019), p. 35.

[4] Roland Pfister et al. "Good things peak in pairs: a note on the bimodality coefficient". In: *Frontiers in Psychology* 4 (2013).

[5] Kleber A Oliveira, Henrique F de Arruda, and Yamir Moreno. "Mechanistic interplay between information spreading and opinion polarization". In: *arXiv preprint arXiv:2410.17151* (2024).

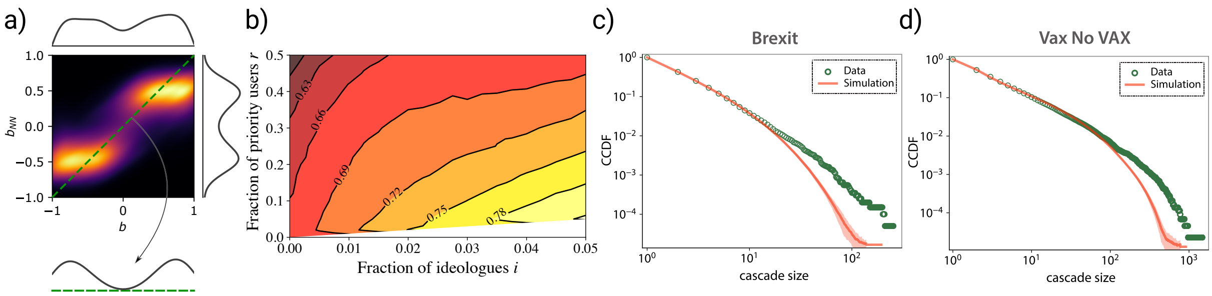

Figure 1: **Main results.** Panel (a) shows the comparison of the polarization and division of the $(b,\, b_{NN})$ with the resulting distribution after rotating the map. Panel (b) displays BC_{hom} by varying the number of ideologues and priority accounts. Panels (c) and (d) show comparisons of the sizes of cascades obtained from Twitter and the simulation after optimization.