Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Combining opinion and structural similarity in link recommendations to counter extreme polarization

Gabriella Dantas Franco

g.dantasfranco@uva.nl

University of Amsterdam

Marta C. Couto

University of Amsterdam

Vítor V. Vasconcelos

University of Amsterdam

Fernando P. Santos

University of Amsterdam

Research Article

Keywords: Complex systems, recommender algorithm, co-evolving network and opinions

Posted Date: July 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9482178/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Applied Network Science on November 6th, 2026. See the published version at https://doi.org/10.1007/s41109-026-00845-9.

Combining opinion and structural similarity in link recommendations to counter extreme polarization

Gabriella Dantas Franco[1*], Marta C. Couto[1],
Vítor V. Vasconcelos[1], Fernando P. Santos[1]

[1]Informatics Institute, University of Amsterdam, Amsterdam, 1098 XH, Netherlands.

*Corresponding author(s). E-mail(s): g.dantasfranco@uva.nl;
Contributing authors: m.gomesdacunhacouto@uva.nl;
v.v.vasconcelos@uva.nl; f.p.santos@uva.nl;

Abstract

Recommendation algorithms, used in online social networks, shape interactions between users. In particular, link-recommendation algorithms suggest new connections and affect how individuals interact and exchange information. These algorithms' efficacy relies on key mechanisms governing the creation of social ties, such as triadic closure and homophily. The first is achieved through *structural* similarity and represents a heightened chance of recommending users to one another given mutual friends; the second is related to *opinion* similarity and conveys an increased chance of recommending a connection given similar individual characteristics. These two mechanisms jointly shape the evolution of social networks and behaviors unfolding over them. Their combined effect on the co-evolution of opinion and structure dynamics remains, however, poorly understood. Here, we study how social networks and opinions co-evolve given the joint effect of rewiring based on opinion and structural similarity. We show that both similarity metrics lead to polarized states, but differ in how they impact network fragmentation and opinion diversity. While strongly relying on opinion similarity leads to a higher variation of opinion, rewiring via network similarity leads to a larger number of (dis)connected components, resulting in fragmented networks that lean towards one of the signed opinions. Under strong homophilic settings, introducing a weak dependence on structural similarity prevents network fragmentation and favors moderate opinions. This work can inform the design of new recommender algorithms that explicitly account for interacting social

and recommendation mechanisms, with the potential to foster moderate opinion coexistence even in inherently polarizing settings.

Keywords: Complex systems, recommender algorithm, co-evolving network and opinions

1 Introduction

Online social networks (OSNs) mediate a large fraction of our social and professional interactions, with various implications for people's offline lives. OSNs affect political mobilization (Bond et al. 2012; Jones et al. 2017), health and well-being (Hobbs et al. 2016; Burke and Kraut 2016), or even access to social and financial capital (Chetty et al. 2022a,b; Steinert-Threlkeld et al. 2015; Greijdanus et al. 2020; Rajkumar et al. 2022). In online environments, new connections are increasingly shaped by algorithmic recommendations. In particular, link recommendation algorithms are used to recommend new connections to users, influencing who connects with whom, how information diffuses, and ultimately how the structure of OSNs evolves (Daly et al. 2010; Su et al. 2016, 2020; Santos et al. 2021).

In order to provide recommendations that users accept, link recommendation algorithms rely, explicitly or implicitly, on well-known mechanisms for social ties formation. *Homophily* and *triadic closure* are two key examples. These mechanisms have long been studied in network science as key principles for the formation of social bonds (Louch 2000). Similarly, homophily and triadic closure suggest important heuristics in link-prediction algorithms (Liben-Nowell and Kleinberg 2003; Li et al. 2017) and are implicitly leveraged by more complex graph representation learning methods (Grover and Leskovec 2016; Kipf and Welling 2016). Public information offered by large platforms to explain algorithmic link recommendations commonly identifies shared connections and common interest as key metrics used, which correspond closely to triadic closure and homophily in network-science terms (Meta 2026; X 2026; LinkedIn

2026). Despite capturing different aspects of social connections and being jointly used in link recommendations, the combined effects of homophily and triadic closure remain poorly understood.

Homophily is an umbrella term used to describe the tendency of individuals to interact with similar others (Lazarsfeld et al. 1954; McPherson et al. 2001). When designing link-prediction or link-recommendation methods, homophily implies that links should be suggested based on *opinion similarity*, i.e., a similarity function comparing nodes based on their opinions, preferences, or attitudes. This metric operates on individuals' properties, accounting for user characteristics that are inferred from their online interactions (Aiello et al. 2012; Sajjadi et al. 2024). Triadic closure, on the other hand, assumes that two disconnected individuals are likely to become linked if they share connections (Peixoto 2022; Bachmann et al. 2025). In general, triadic closure is based on a similarity function that assumes that two nodes are similar if they share a high number of neighbors. Such similarity function is often called *structural similarity* as it only depends on the social network structure (Lü and Zhou 2011) – in opposition to opinion similarity, which depends on individuals' preferences and opinions.

Understanding how rewiring based on opinion and structural similarity shape network structure and opinion dissent is therefore essential for assessing the long-term effects of link recommendation systems. Disentangling homophily and triadic closure as causal mechanisms in empirical data has only been addressed recently by Peixoto (2022). Moreover, their cumulative effects have been investigated by Asikainen et al. (2020) on a population with static opinions. More theoretically, Abebe et al. (2022) have studied the effect on network segregation when adding an edge via triadic closure on stochastic block models with homophily. Bachmann et al. (2025) studied the joint effects of triadic closure and homophily in inequality between groups. In all previous works mentioned, homophily depends on individual characteristics and group membership that remain static over time. In another direction, previous work studied

coevolution of node properties and network topology (Gross and Blasius 2008) and it has been shown that, when opinions co-evolve with homophilic rewiring, there is a continuous phase transition in structure Holme and Newman (2006). However, we still lack understanding of the joint effects of homophily and triadic closure in recommender systems, particularly in domains where both the network and opinions – or, broadly, nodes' attributes – co-evolve. We are compelled to understand the co-evolution of node attributes and network structure as this setup can capture ubiquitous phenomena in OSNs such as feedback loops, the emergence of echo-chambers and the emergence of polarized social structures detrimental to collective action (Vasconcelos et al. 2021). Disentangling peer influence effects from homophily and triadic closure, all of which lead to dense networks of like-minded individuals, is particularly hard empirically due to confounded mechanisms and various feedback loops (Su et al. 2020).

To address this gap, we propose a model in which opinions and network structure co-evolve. Opinions are updated based on neighbors average influence following an update rule recently proposed as capturing the formation of echo chambers in online social media platforms (Baumann et al. 2020). Link rewiring, on the other hand, is driven by a tunable combination of opinion and structural similarity. By tuning both the strength and interpolation between opinion and structural similarity, we aim at testing mechanistically different forms of link-recommendation applied in OSNs. The co-evolving dynamics of opinions and links leads to distinct regimes of polarization and network fragmentation, highlighting the importance of algorithmic design choices. We ask whether algorithms that prioritize structural similarity over opinion similarity lead to qualitatively different patterns of network connectivity and whether certain parameter regimes sustain opinion coexistence without extreme polarization and network segregation.

Our results reveal differences between structural and opinion-driven rewiring. Both lead to polarized and fragmented communities that hold opposing opinions and are

not connected to each other. However, opinion-based rewiring leads to a fragmented state more easily, with a lower threshold for the fragmentation transition. We also show that a weak reliance on structural-based triadic closure can preserve a connected network, allowing for the coexistence of moderate opinions within the same interacting structure. This highlights that a small amount of structural similarity might be needed to counteract the inherently polarizing aspect of homophily.

2 Methods

To investigate the effects of link-recommendation algorithms on population-level opinion and network structure, we develop a mechanistic co-evolutionary model that accounts for rewiring based on opinion and structural similarity metrics. We aim to understand their impact on long-term outcomes in the user population. In this section, we detail the models used in this study. First, we present the recommender design, which interpolates between opinion and structural similarity. Secondly, we introduce the opinion dynamics model. Lastly, we discuss the metrics used to capture opinion dynamics.

2.1 Network dynamics through link-recommendations

Link-recommendation algorithms used in OSNs involve complex algorithms and pipelines (Meta 2026; LinkedIn 2026; Twitter, Inc. 2023; Gupta et al. 2013). Fundamentally, however, these systems often rely on a combination of users' activity, interests and network data. Here, we attempt to isolate, in a stylized manner, two key metrics used in link-recommendations (explicitly as heuristics or implicitly in complex graph learning methods): opinion similarity and structural similarity. We assume that nodes will be ranked, and and top candidates are offered as recommendations to users. The selection process is simulated by a random draw from the probability distribution of recommended not-yet-friends, and we assume that, once drawn, the recommended user

is accepted as a friend, establishing the new connection. Once a new link is formed, the focal individual will choose one of its previously existing connections to break, with uniform probability. This assumption is based on the 'attention-budget phenomena', an expression commonly used to state that despite the increasing number of connections on OSNs, each individual is only able to reciprocate attention with a few other users (Huberman et al. 2008; Weng et al. 2012). It also keeps the number of total edges constant over time, which facilitates analytic approaches.

The population is represented by an undirected graph G, where each node u is a user ($u \in V$), and each link l is an active connection ($l \in E$). E denotes the set of edges (links) and V the set of vertices (nodes/users). The network changes are captured by a dynamic adjacency matrix A_{ij}^t, where $A_{ij}^t = 1$ if an edge between node i and node j exists at time t, and $A_{ij}^t = 0$ otherwise. To interpolate between recommendations based on opinion similarity and structural similarity, we assume that the probability of adding a new link between nodes i and j is given by

$$P[A_{ij}^{t+1} = 1] = \rho\, H_{ij}(t) + (1 - \rho)\, S_{ij}(t), \tag{1}$$

where H_{ij} and S_{ij} represent the similarity measures based on opinion and structure, respectively. The parameter ρ can be interpreted as the weight of opinion similarity on recommendations, in opposition to structural similarity. At one extreme, $\rho = 0$ represents a rewiring dynamics that is purely guided by the structural similarity metric, recovering (Santos et al. 2021). On the other extreme, $\rho = 1$ considers only opinion similarity, and does not take structure into account, approximating models where homophily shapes the rewiring mechanism (Holme and Newman 2006; Blex and Yasseri 2022; Borges et al. 2024; Vasconcelos et al. 2019; Gross and Blasius 2008).

At each time step, each node i will have a turn as focal, selecting a new node j to add to its neighborhood set N_i. Once selected, the friendship is accepted with probability

1, leading to $A_{ij}^{t+1} = 1$. The new friendship connection is drawn from the focal node's recommendations, given by the probability distribution $P[A_{ij}^{t+1} = 1]$ defined above.

2.1.1 Structural similarity

For the purpose of this study, structural similarity between nodes i and j is represented by the number of common neighbors, or the cardinality of the intersection between the neighborhood sets of i and j, $|N_i \cap N_j|$. Although simple, this metric has been used extensively in link-prediction methods with great success (Liben-Nowell and Kleinberg 2003). Common neighbors statistics have been estimated from empirical network data in frameworks such as Stochastic Actor-Oriented Models (Snijders 2001; Goodreau et al. 2009), lending plausibility to its use as an empirically testable mechanism. Moreover, to allow for links between different groups to be formed, we include a fixed amount of noise ϵ, that can also be interpreted as algorithmic uncertainty. A non-zero noise is also needed to avoid divergences on the similarity metrics. For all nodes j not currently connected to the focal i, we then have that

$$S_{ij}(t) = \frac{\left[|N_i^t \cap N_j^t|(1-2\epsilon) + \epsilon\right]^\eta}{\sum_{k \notin N_i} \left[|N_i^t \cap N_k^t|(1-2\epsilon) + \epsilon\right]^\eta},$$

where $\eta(\geq 0)$ controls the strength of the similarity metric. The higher this exponent, the more skewed the distribution is.

2.1.2 Opinion similarity

We construct the opinion similarity metric similarly, but taking into consideration the opinion space instead of the network structure. As detailed below, we assume that each individual has a real-valued opinion $x_i \in \mathbb{R}$ over a single, isolated topic. Given this one-dimensional opinion space, a natural proxy for homophily is the distance between individual opinions – subject to the same noise ϵ. Analogously to the structural similarity function, there is a non-linear parameter $\beta(\geq 0)$ representing the strength

of similarity, leading to

$$H_{ij}(t) = \frac{\left[|x_i^t - x_j^t|(1-2\epsilon)+\epsilon\right]^{-\beta}}{\sum_{k\notin N_i}\left[|x_i^t - x_k^t|(1-2\epsilon)+\epsilon\right]^{-\beta}},$$

where now $|x|$ denotes the absolute value of x.

Once again, the metric is normalized for each focal node i, so that each user selects exactly one candidate from its distribution. The negative sign on the exponent turns the distance between opinions, which measures how dissimilar the node opinions are, into a similarity metric.

This setup fully describes the recommender with only a few parameters: the exponents η and β independently tune the strengths of structural and opinion similarity. Importantly, these parameters allow us to test different dependencies of recommendations on similarity metrics: a sub-linear ($0 < \eta < 1$), linear ($\beta, \eta = 1$), and supra-linear ($\beta, \eta > 1$) dependency, later shown to be important. Additionally, $\rho \in [0, 1]$ balances the weight of opinion similarity against structural similarity. And ϵ defines algorithmic noise, which is considered constant. For some more discussion over ϵ and its effect on this dynamics, please refer to the Supplementary Information.

2.2 Opinion Dynamics

As introduced above, we aim at understanding the effects of link recommendations on a population with evolving opinions. We adopt the discrete-time counterpart of a model recently identified as suitable to model echo chambers in online social networks (Baumann et al. 2020). This model assumes that opinions change given the average opinion on a node's neighborhood; under random rewiring, this type of local averaging dynamics leads to consensus being an absorbing state (DeGroot 1974).

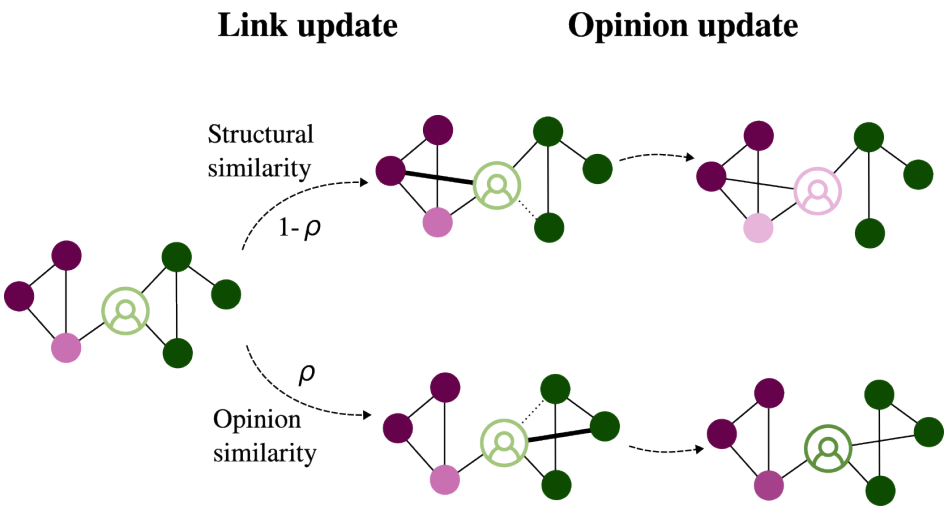


Fig. 1 Illustration of our model. In this example, we assume that the center node (light green) will rewire one of its connections. This node receives a friend recommendation based on opinion similarity (ρ) or structural similarity ($1-\rho$) and accepts it (thick edge). At the same time, it chooses a random existing link to exclude (dotted edge). Once the whole network has had a chance to rewire, all nodes update their opinion simultaneously, as indicated by the color changes, where darker colors indicate more extreme opinions.

Each subject i has an individual opinion $x_i^t \in \mathbb{R}$ about a single topic. Social influence is modeled as a synchronous process: once every user has had the opportunity to rewire one of its connections, all of them revisit their opinions at the same time, influenced by their respective neighborhoods. As introduced above, we consider the average opinion of a node's neighborhood; this effect is however non-linear: users with extreme opinions have a bounded amount of influence over their peers. This assumption is supported by empirical evidence (Jayles et al. 2017) and has been studied in the context of online activity-driven networks (Baumann et al. 2020, 2021).

In the absence of peer influence, individuals forget the topic, and their opinion decays towards a neutral stance at a constant rate $0 \leq \gamma < 1$. This can be interpreted as referring to a topic that is not a core characteristic of the individuals: it is only important when others are talking about it, much like 'hot topics'.

Building on these observations, formally, the opinion dynamics we consider is given by

$$x_i^{t+1} = \gamma\, x_i^t + K \sum_{j=1}^{N} \frac{A_{ij}^t}{k_i^t} \tanh(\alpha\, x_j^t), \tag{2}$$

where $k_i^t = \sum_l A_{il}^t$ is the degree of node i and K controls social influence. Lastly, α controls how one's opinion translates into social influence, representing the interplay between issue controversy and social reinforcement of this model (Baumann et al. 2020; Santos et al. 2021). We illustrate the whole model in Fig. 1.

On a static network (or under random rewiring), low α leads to a stable solution of neutral consensus ($x_i = 0,\ \forall i$); for higher values, persistent non-neutral stances emerge (Gray et al. 2018; Bizyaeva et al. 2022). Although we allow any real value for the opinion variable, the balance between opinion decay and social influence effectively bounds the opinions in the range $[\frac{-K}{1-\gamma}, \frac{K}{1-\gamma}]$.

Unless otherwise stated, we set the parameters $K = 0.1$, $\gamma = 0.99$ and $\alpha = 0.3$ (see (Baumann et al. 2020, 2021; Santos et al. 2021) for more details on this model's behavior and interpretation on online social settings). Under these values, any connected static or randomly rewired network converges to a non-neutral but consensual state. When connections adapt through the recommender system, however, this long-term outcome may change – so that *any polarization observed, either in opinion or in network structure (e.g., fragmentation), arises solely from the network dynamics and, to a large extent, the link-recommendation algorithm assumed.*

2.3 Metrics

We use three main metrics to interpret our results: polarization, radicalization and the number of connected components. First, we are interested in characterizing a real-valued stance within a population, so it is sensible to choose a dispersion (or variance-)based

measurement for the opinions. For that purpose, we define **polarization** as

$$\text{Polarization}(\mathbf{X}) \equiv \text{std}(\mathbf{X}),$$

i.e., the standard deviation of the opinion vector $\mathbf{X} = (x_1, x_2, \ldots, x_N)$. As a dispersion-based measure, this definition of polarization is a good representation of dissent in a population. Polarization is only null when there is complete consensus.

We are also interested in how extreme or radical the opinions can be. Previous work considers the mean of absolute opinions as the radicalization level within a population (Santos et al. 2021). We follow a similar approach and define **radicalization** as

$$\text{R}(\mathbf{X}) \equiv \frac{1}{N} \sum_{i=1}^{N} |x_i|.$$

Polarization is a variance-based metric that can take into account both in-group dispersion and between-group dispersion. However, many distinct network structures can display similar opinion distributions. As such, this metric does not capture *structural polarization*, as described in (Borges et al. 2024). There, besides attribute-level dissent, there is also a network-level separation, or fragmentation. Therefore, to account for fragmentation, we also consider the **number of connected components** of the network. A connected component of G is a maximal connected subset of vertices $C \subseteq V$, i.e., for all $u, v \in C$ there exists a path between u and v, and no vertex in $V \setminus C$ is connected by a path to any vertex of C.

3 Results

Recall that the effect of opinion and structural similarity-based metrics on recommendations is controlled by two parameters (β and η, respectively). Furthermore, the relative importance of opinion similarity over structural similarity is controlled by ρ.

Initially, we assume that $\eta = \beta$, allowing us to analyze the emerging levels of polarization for the 2-dimensional parameter space ($\eta = \beta,\ \rho$) shown in Fig. 2. First, we note that, for any value of ρ – that is, whether the re-linking prioritizes the opinion or structural similarity – there exists a threshold on similarity strength (β, η) after which polarization emerges. This suggests that both similarity mechanisms can contribute to polarization. However, the highest values of polarization ensue when recommendations prioritize opinion similarity (high ρ), a setting where we observe a sudden transition between consensus and polarization.

The bottom panels of Fig. 2 illustrate typical opinion timeseries for the highlighted parameter regions. In region I (leftmost), below the critical value for the similarity strength, opinion consensus emerges. Under this parameter regime, any random uniform initial condition converges with equal chance to the extreme negative or positive stance, similarly to random recommendations. In region II (upper-right), which corresponds to opinion similarity being favored in recommendations ($\rho > 0.5$), the average opinion over the population remains close to null. The transition from consensus to polarized states occurs near the linear recommendation ($\eta = \beta = 1$). On the other hand, region III (lower-right) showcases how prioritizing structural similarity in recommendations leads to the average opinion over the population being biased towards one direction – explaining partly why we observe lower polarization when ρ is low. Here, the threshold values of η above which polarization is non-zero are supra-linear. This suggests that consensus persists for a wider range of similarity strengths when recommendations favor structural similarity.

Inspecting polarization results alone might suggest that opinion and structural similarity play qualitatively similar roles in leading to polarized communities, apart from differences in polarization values and response to changes in similarity rewiring strength. However, looking at the number of connected components over the same

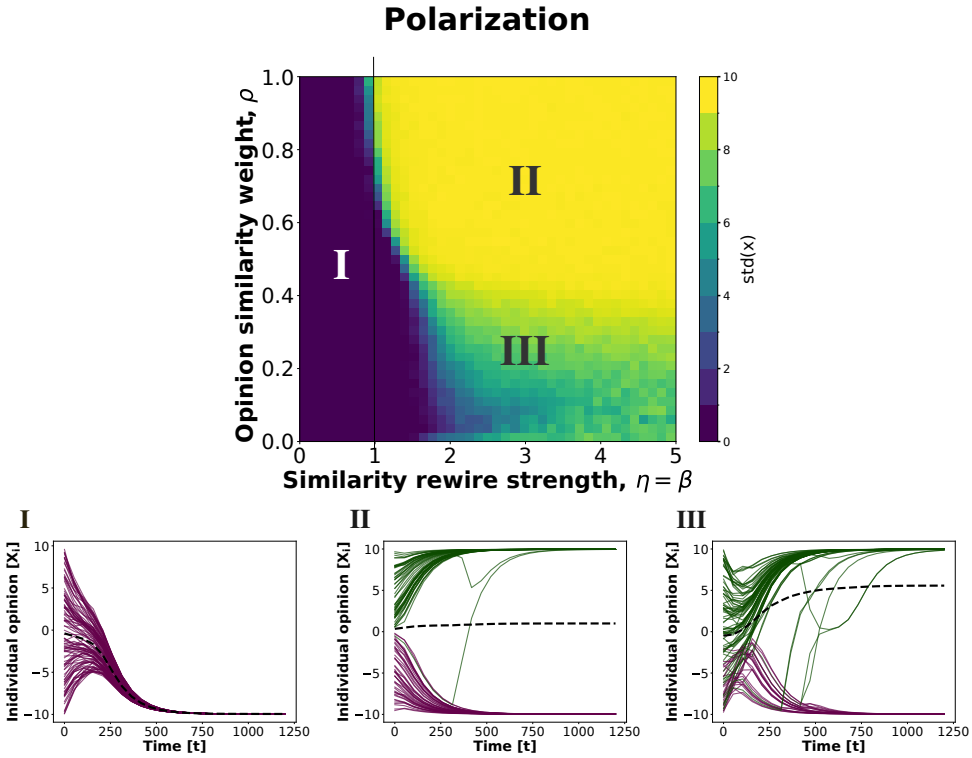


Fig. 2 (Top) Polarization when varying the dependency of recommendations on opinion and structural similarity (x-axis, $\eta = \beta$) and the relative importance of opinion similarity (homophily) on recommendations (y-axis, ρ). Each point corresponds to the average polarization over 50 independent simulations. We observe that polarization emerges for any value of ρ, given a sufficiently high value of β and η (here we assume $\beta = \eta$). This suggests that strong homophily and triadic closure each independently contribute to opinion polarization. (Bottom) Timeseries of opinion evolution for three paradigmatic regimes highlighted in the top panel: I ($\eta = 0.5$, $\rho = 0.5$), II ($\eta = 3.0$, $\rho = 0.75$), III ($\eta = 3.0$, $\rho = 0.25$). Each line represents an agent's opinion over time, colored according to its value at $t = 1200$. Dashed lines indicate the average opinion of the population. We confirm that strongly relying on similarity-based rewiring results in a bi-modal and highly polarized opinion landscape. Other parameters considered: $\alpha = 0.3$, $K = 0.1$, $\gamma = 0.99$, $t_{max} = 1200$.

parameter space, shown in Fig. 3, we see that fragmentation levels differ greatly depending on the recommender's dependency on opinion or structural similarity. For $\rho < 0.5$, the number of connected components is high, having two peaks: a global maximum at $\rho = 0$ and a local one at intermediate values of ρ. For $\rho > 0.5$, the system generally reaches a stable structure of 2 connected components that do not exchange information, which is expected in highly homophilic systems (as investigated by Baumann et al. (2021), for example). We also note that such fragmented regime is

reached with a lower similarity strength when the recommender prioritizes opinion similarity. This emphasizes that, in the case of equal similarity strengths ($\beta = \eta$), the network is more resilient to fragmentation when recommendations rely more on structural similarity.

Examples of the resulting networks in each region are displayed on the bottom panels of Fig. 3. These network examples depict two general trends observed: 1) High similarity strength with the recommender relying on opinion similarity leads to two groups of roughly the same size; and 2) High similarity strength with the recommender relying on structural similarity leads to a fragmented state with a tightly connected majority group, and many small opposing communities.

The previous results indicate that when both η and β are high (i.e., recommendations are highly sensitive on opinion and structural differences between users) interpolating between these two features (changing ρ) leads to significantly different network topologies. This means that, even if grounding link recommendations on either opinion or structural similarity can increase polarization, the distribution and magnitude of opinions can vary as a result of different network configurations. Understanding how to combine high values of η and β to mitigate the negative effects of polarization is particularly relevant in our context as one needs to guarantee that recommendations are still accurate and useful for users (Li et al. 2017), thereby relying on their interests and acquaintances (i.e., $\beta = \eta = 0$ might lead low polarization, however such recommenders will hardly be adopted given their random nature).

Beyond polarization, we now turn our attention to radicalization, a measure of how extreme opinions are. We start by noting that polarization and radicalization can vary independently: even if radicalization is high, polarization depends on opinions' dispersion within opposing signed groups, on the relative group sizes and the exact magnitude of opinion within each group (we formally describe the relationship between

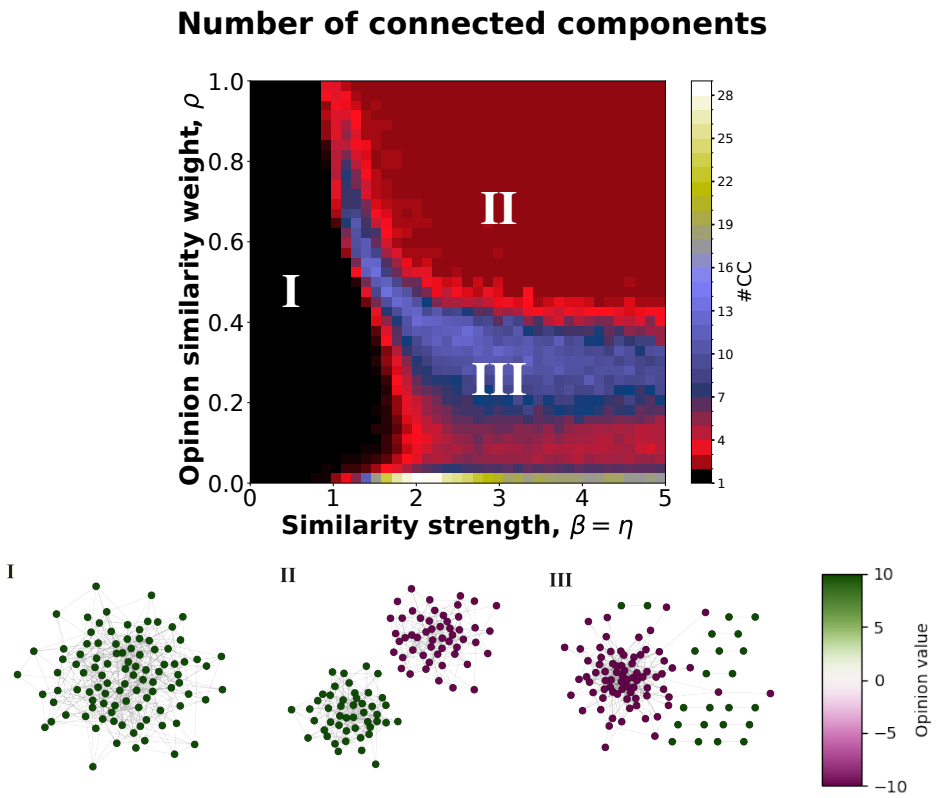


Fig. 3 Strong homophily and transitivity each independently contribute to opinion polarization, however this occurs through the emergence of different network topologies. (Top) Average number of connected components ($\#CC$) over 50 independent simulations. The color scale emphasizes the regions with: (I) a single connected component; (II) approximately two connected components, which corresponds to prioritizing homophily ($\rho > 0.5$); and (III) many connected components, which appear when the recommender prioritizes transitivity ($\rho < 0.5$), showing a high level of fragmentation. (Bottom) Long-term network structures for the three paradigmatic regimes highlighted in the parameter space. Other parameters considered: $\alpha = 0.3$, $K = 0.1$, $\gamma = 0.99$, $t_{max} = 3000$.

polarization and radicalization in the Supplementary Information, Appendix A). We thereby ask: how to set ρ under high values of β in order to reduce radicalization?

To answer this question we allow both similarity strengths η and β to vary independently and focus on varying η under high values of β. In Fig. 4, we note that the level of opinion similarity (ρ) – together with a sublinear value for structural strength ($\eta < 1$) – determines the regime in which different opinions coexist. For low values of homophily ($\rho < 0.5$), we observe dominance of one single stance over the other, and the population reaches consensus. For intermediate values of homophily ($0.5 < \rho < 0.75$),

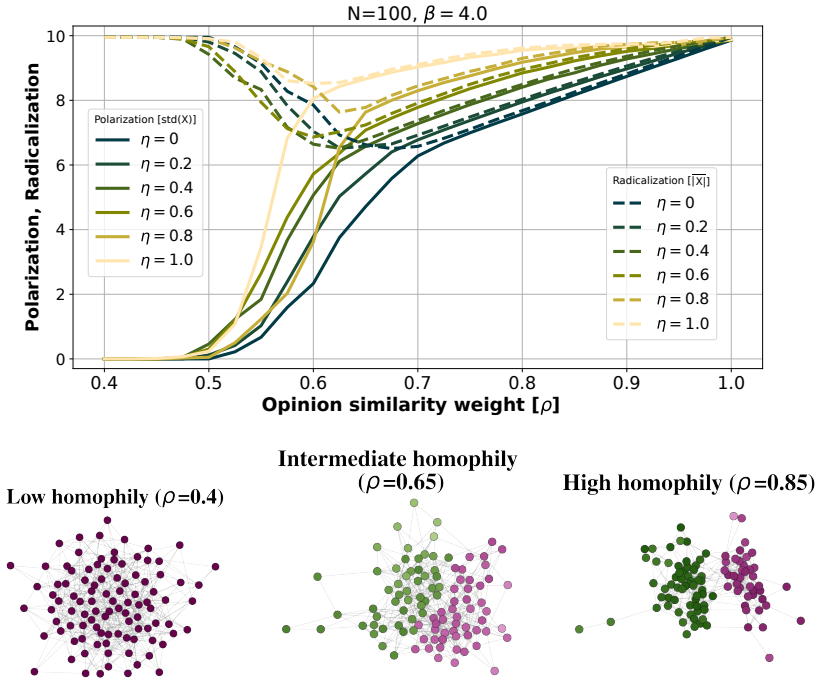


Fig. 4 (Top) Polarization (solid lines) and radicalization (dashed lines) values for $\rho > 0.4$. While polarization rises for all 6 values of η shown, sublinear values are more efficient in dampening extremism, with lowest radicalization values found for $\eta < 0.4$. As seen by the convex aspect of the radicalization curves, the minimum appears at intermediate values of ρ that depend on other parameters. (Bottom) Network representations illustrating three different regimes observed. For low values of ρ the consensus state is reached. For intermediate levels, we observe a global minimum on the radicalization curve, accounting for less extremists on the connected population. For high values of ρ we observe an increase in radicalization, led by an increasing separation between groups.

we observe a minimum value for radicalization for that set of parameters. The corresponding network does not clearly favor one of the opinions over the other, as can be seen by two similarly sized modules with strong internal connectivity. On the other hand, for high values of opinion similarity ($\rho > 0.75$), radicalization values start to rise, together with polarization, as we see individual opinions becoming more extreme (darker colors) and the modules becoming less connected to each other. We show in the Supplementary Information (also for $\beta = 4$) that useful recommendations through triadic closure (positive η) are only able to lower radicalization to a minimum until a certain weight of homophily on the recommender.

As briefly discussed above, we see that strong homophily ($\rho \approx 1$) leads to fragmentation under linear similarity strength ($\eta = 1$). On the other hand, favoring transitivity ($\rho < 0.5$) allows for a single connected component until supra-linear conditions ($\eta > 1$), resulting in a system that is more resilient against fragmentation. Together with our observations from Fig. 4, this indicates that a certain level of transitivity is needed to prevent the population from breaking apart, while possibly keeping coexistence of opinions.

Although the results we discuss here are numerical and cover a limited range of network sizes, in the Supplementary Information we provide a robustness check for more parameters and an analytic study of the interplay between opinion and structural similarity metrics in preventing network fragmentation. Such analysis qualitatively confirms that, especially on homophilic settings, a certain level of transitivity is needed to counteract the tendency for group-wise fragmentation (see Supplementary Information, Appendix B).

4 Conclusion

Here we present a new model that considers both structural similarity (based on triadic closure) and opinion similarity (based on homophily) when recommending and adding links in online social networks. We aim at studying how these two distinct principles of users' similarity, when combined in link-recommendation algorithms, affect opinion dynamics and network structure. While homophily and triadic closure are well-known mechanisms for creating social ties, their combined effects remain largely unexplored, particularly in settings where users' preferences co-evolve with network structure. Previous work (Abebe et al. 2022; Bachmann et al. 2025; Asikainen et al. 2020) investigated the impact of triadic closure on network formation models that consider homophily. Under these static-feature assumptions, observations are not straightforward. While (Abebe et al. 2022; Bachmann et al. 2025) have shown that

triadic closure decreases network segregation, hinting that the two principles might not work in tandem, (Asikainen et al. 2020) observes an increase in observed homophily.

Our results show that, despite inherent differences in the similarity metrics, both can independently lead to polarized communities when link-recommendation algorithms strongly rely on them as heuristics to suggest links. We also show that polarization characteristics differ among similarity metrics: while recommendations based on opinion similarity lead to higher polarization, rewiring based on structural similarity leads to a larger number of (dis)connected components. Nevertheless, we demonstrate that, under strong homophilic scenarios, rooting recommendations on structural similarity can provide a route to offer useful recommendations to users, while dampening extremism and enabling opinion coexistence. The above observations agree with Abebe et al. (2022), as they highlight that a weak but existent relationship with transitivity appears to be crucial for avoiding network fragmentation. Additionally, a recent work by Bachmann et al. (2025) argues that when triadic closure is conditional on the mechanisms of homophily and preferential attachment it exacerbates degree inequality while reducing segregation. Taken together, these results suggest that, although polarization emerges as a robust outcome of similarity-driven recommendation, incorporating structural similarity as a principle for recommendations can counteract the high polarization and radicalization outcomes resulting from high homophily.

Our findings are conditioned on two modeling choices: an opinion dynamics model that only results in polarization when due to the recommender, and simplifying assumptions of the network (fixed size, density and initial topology). Exploring other classes of opinion dynamics on our model would be a natural pathway for testing robustness of rewiring-driven polarization. For example, bounded confidence models could identify tolerance thresholds that alter the observed stable states; or anti-conformist individuals could act opposed to the averaging rule, changing the absorbing states of the dynamics (as in Santos et al. (2021) and Mittal et al. (2024)).

More broadly, our results show that design choices on online social platforms' recommender systems can deeply affect polarization and network cohesion in the long-term. When models are universally applied at a given platform, design and parameter choices can have real large-scale impacts. Here, we show that, counter-intuitively, transitivity appears to compensate for homophilic tendencies. This observation by itself incentivizes further research to investigate the matter, and leads a new way for online recommenders to be understood and designed for social media platforms.

Declarations

Availability of data and materials

Data sharing is not applicable to this article as no datasets were generated or analyzed during the current study.

Competing interests

The authors have no conflicts of interest to declare that are relevant to the content of this article.

Funding

This work is supported by ERC grant (RE-LINK, https://doi.org/10.3030/101116987). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Authors' contributions

G.D.F. and F.P.S. contributed to the conceptualization and research design of this work. G.D.F. was involved in coding, implementing the experiments, analyzing the

results and writing. F.P.S., M.C.C. and V.V.V. were involved in supervision and revision of the final draft. All authors read and approved the final manuscript.

References

Abebe R, Immorlica N, Kleinberg J, et al (2022) On the effect of triadic closure on network segregation. In: Proceedings of the 23rd ACM Conference on Economics and Computation, pp 249–284

Aiello LM, Barrat A, Schifanella R, et al (2012) Friendship prediction and homophily in social media. ACM Transactions on the Web (TWEB) 6(2):1–33

Asikainen A, Iñiguez G, Ureña-Carrión J, et al (2020) Cumulative effects of triadic closure and homophily in social networks. Science Advances 6(19):eaax7310

Bachmann J, Martin-Gutierrez S, Espín-Noboa L, et al (2025) Network inequality through preferential attachment, triadic closure, and homophily. arXiv preprint arXiv:250923205

Baumann F, Lorenz-Spreen P, Sokolov IM, et al (2020) Modeling echo chambers and polarization dynamics in social networks. Physical Review Letters 124(4):048301

Baumann F, Lorenz-Spreen P, Sokolov IM, et al (2021) Emergence of polarized ideological opinions in multidimensional topic spaces. Physical Review X 11(1):011012

Bizyaeva A, Franci A, Leonard NE (2022) Nonlinear opinion dynamics with tunable sensitivity. IEEE Transactions on Automatic Control 68(3):1415–1430

Blex C, Yasseri T (2022) Positive algorithmic bias cannot stop fragmentation in homophilic networks. The Journal of Mathematical Sociology 46(1):80–97

Bond RM, Fariss CJ, Jones JJ, et al (2012) A 61-million-person experiment in social influence and political mobilization. Nature 489(7415):295–298

Borges HM, Vasconcelos VV, Pinheiro FL (2024) How social rewiring preferences bridge polarized communities. Chaos, Solitons & Fractals 180:114594

Burke M, Kraut RE (2016) The relationship between facebook use and well-being depends on communication type and tie strength. Journal of Computer-Mediated Communication 21(4):265–281

Chetty R, Jackson MO, Kuchler T, et al (2022a) Social capital i: measurement and associations with economic mobility. Nature 608(7921):108–121

Chetty R, Jackson MO, Kuchler T, et al (2022b) Social capital ii: determinants of economic connectedness. Nature 608(7921):122–134

Daly EM, Geyer W, Millen DR (2010) The network effects of recommending social connections. In: Proceedings of the Fourth ACM Conference on Recommender Systems, pp 301–304

DeGroot MH (1974) Reaching a consensus. Journal of the American Statistical association 69(345):118–121

Goodreau SM, Kitts JA, Morris M (2009) Birds of a feather, or friend of a friend? using exponential random graph models to investigate adolescent social networks. Demography 46(1):103–125

Gray R, Franci A, Srivastava V, et al (2018) Multiagent decision-making dynamics inspired by honeybees. IEEE Transactions on Control of Network Systems 5(2):793–806

Greijdanus H, de Matos Fernandes CA, Turner-Zwinkels F, et al (2020) The psychology of online activism and social movements: Relations between online and offline collective action. Current Opinion in Psychology 35:49–54

Gross T, Blasius B (2008) Adaptive coevolutionary networks: a review. Journal of the Royal Society Interface 5(20):259–271

Grover A, Leskovec J (2016) node2vec: Scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp 855–864

Gupta P, Goel A, Lin J, et al (2013) Wtf: The who to follow service at twitter. In: Proceedings of the 22nd International Conference on World Wide Web, pp 505–514

Hobbs WR, Burke M, Christakis NA, et al (2016) Online social integration is associated with reduced mortality risk. Proceedings of the National Academy of Sciences 113(46):12980–12984

Holme P, Newman ME (2006) Nonequilibrium phase transition in the coevolution of networks and opinions. Physical Review E—Statistical, Nonlinear, and Soft Matter Physics 74(5):056108

Huberman BA, Romero DM, Wu F (2008) Social networks that matter: Twitter under the microscope. arXiv preprint arXiv:08121045

Jayles B, Kim Hr, Escobedo R, et al (2017) How social information can improve estimation accuracy in human groups. Proceedings of the National Academy of Sciences 114(47):12620–12625

Jones JJ, Bond RM, Bakshy E, et al (2017) Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 us presidential election.

PLOS ONE 12(4):e0173851

Kipf TN, Welling M (2016) Variational graph auto-encoders. arXiv preprint arXiv:161107308

Lazarsfeld PF, Merton RK, et al (1954) Friendship as a social process: A substantive and methodological analysis. Freedom and Control in Modern Society 18(1):18–66

Li Z, Fang X, Sheng ORL (2017) A survey of link recommendation for social networks: Methods, theoretical foundations, and future research directions. ACM Transactions on Management Information Systems (TMIS) 9(1):1–26

Liben-Nowell D, Kleinberg J (2003) The link prediction problem for social networks. In: Proceedings of the Twelfth International Conference on Information and Knowledge Management, pp 556–559

LinkedIn (2026) Linkedin help center. URL https://www.linkedin.com/help/linkedin/answer/a544682/people-you-may-know-feature-

Louch H (2000) Personal network integration: transitivity and homophily in strong-tie relations. Social Networks 22(1):45–64

Lü L, Zhou T (2011) Link prediction in complex networks: A survey. Physica A: Statistical Mechanics and Its Applications 390(6):1150–1170

McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: Homophily in social networks. Annual Review of Sociology 27(1):415–444

Meta (2026) Meta transparency center. URL https://transparency.meta.com/features/explaining-ranking/fb-people-you-may-know/

Mittal D, Constantino SM, Vasconcelos VV (2024) Anticonformists catalyze societal transitions and facilitate the expression of evolving preferences. PNAS Nexus 3(8):pgae302

Peixoto TP (2022) Disentangling homophily, community structure, and triadic closure in networks. Physical Review X 12(1):011004

Rajkumar K, Saint-Jacques G, Bojinov I, et al (2022) A causal test of the strength of weak ties. Science 377(6612):1304–1310

Sajjadi S, Martin-Gutierrez S, Karimi F (2024) Unveiling homophily beyond the pool of opportunities. arXiv preprint arXiv:240113642

Santos FP, Lelkes Y, Levin SA (2021) Link recommendation algorithms and dynamics of polarization in online social networks. Proceedings of the National Academy of Sciences 118(50):e2102141118

Snijders TA (2001) The statistical evaluation of social network dynamics. Sociological Methodology 31(1):361–395

Steinert-Threlkeld ZC, Mocanu D, Vespignani A, et al (2015) Online social networks and offline protest. EPJ Data Science 4(1):1–9

Su J, Sharma A, Goel S (2016) The effect of recommendations on network structure. In: Proceedings of the 25th International Conference on World Wide Web, pp 1157–1167

Su J, Kamath K, Sharma A, et al (2020) An experimental study of structural diversity in social networks. In: Proceedings of the International AAAI Conference on Web and Social Media, pp 661–670

Twitter, Inc. (2023) Follow recommendations service. https://github.com/twitter/the-algorithm/blob/main/follow-recommendations-service/README.md, gitHub

repository README, accessed 2026-03-11

Vasconcelos VV, Levin SA, Pinheiro FL (2019) Consensus and polarization in competing complex contagion processes. Journal of the Royal Society Interface 16(155):20190196

Vasconcelos VV, Constantino SM, Dannenberg A, et al (2021) Segregation and clustering of preferences erode socially beneficial coordination. Proceedings of the National Academy of Sciences 118(50):e2102153118

Weng L, Flammini A, Vespignani A, et al (2012) Competition among memes in a world with limited attention. Scientific Reports 2(1):335

X (2026) X help center. URL https://help.x.com/en/resources/recommender-systems/account-recommendations

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Supplementary.pdf