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.

Polarization and echo chambers in Reddit's political discourse

Daniele Cirulli
University of Rome "Tor Vergata"
Antonio Desiderio
antonio.desiderio@isi.it
Institute for Scientific Interchange
Giulio Cimini
University of Rome "Tor Vergata"
Fabio Saracco
Enrico Fermi Center for Study and Research

Research Article

Keywords: Social Networks, Reddit, Statistical Validation, Computational Social Science

Posted Date: June 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9813312/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 EPJ Data Science on October 12th, 2026.
See the published version at https://doi.org/10.1140/epjds/s13688-026-00707-5.

Polarization and echo chambers in Reddit's political discourse

Daniele Cirulli[1,2], Antonio Desiderio[1,2,3,4], Giulio Cimini[*,2], Fabio Saracco[*,2,5,6]

[1]Physics Department and INFN, University of Rome "Tor Vergata", Via della Ricerca Scientifica 1, 00133 Rome, Italy.
[2]"Enrico Fermi" Research Center, Via Panisperna 89A, 00184 Rome, Italy.
[3]ISI Foundation, Via Chisola 5, 10126 Turin, Italy.
[4]Department of Applied Mathematics and Computer Science, Technical University of Denmark, Richard Petersens Plads, 2800 Lyngby, Denmark.
[5]Institute for Applied Computing "Mauro Picone", National Research Council of Italy, Via Madonna del Piano 10, 50019 Sesto Fiorentino, Italy.
[6]IMT School For Advanced Studies Lucca, Piazza S.Francesco 19, 55100 Lucca, Italy.

Abstract

Political debate nowadays takes place mainly on online social media, with election periods amplifying ideological engagement. Reddit is generally considered more resistant to polarization and echo chamber effects than platforms like Twitter or Facebook. Here, we challenge this assumption through a case study across the 2016 US presidential election. We use statistical validation techniques to extract ideologically distinct communities of subreddits, in terms of their contributing user base and news consumption, which we use to analyze the dynamics of political debate. We thus reveal clear polarization in both interaction-based and topic-based communities, with clusters of Democratic, Conservative, and Banned subreddits. Election periods intensify cross-group engagement, align Banned and Conservative content, and reduce linguistic diversity within groups. Overall we characterize Reddit as a polarized environment marked by the presence of echo chambers, highlighting network validation as a key method for identifying behavioral and interaction patterns on online social media.

Keywords: Social Networks, Reddit, Statistical Validation, Computational Social Science

Introduction

Digital platforms have become increasingly central in shaping political discourse [1–3]. Political campaigns are often tested online and, if successful, promoted through legacy mass media or covered as native digital events in their own right [4]. This growing societal tendency to communicate in online settings offers researchers unprecedented access to quantitative traces [5], through which social behavior can be studied on different scales [6]. Due to early data availability, their popularity and centrality in political discourse, much of the literature has focused on *Facebook* and *X* (formerly *Twitter*). Various studies have shown that political discussion on these platforms often produces tightly clustered user communities, called echo chambers, in which like-minded individuals reinforce their prior beliefs, often in opposition to external views [7–14]. Whether such patterns generalize to other platforms, potentially characterized by different user experiences and interaction dynamics, remains an open question [15, 16].

For instance, studies on *Reddit* report no strong presence of echo chambers [12, 17, 18], possibly depending on the resolution scale of the analysis [19, 20], despite potentially high levels of polarization [21]. This outcome is particularly surprising, as the structural design of Reddit is expected to facilitate, rather than hinder, the formation of echo chambers. Indeed, unlike timeline-based platforms, Reddit is structured into thematic *subreddits* – user-defined forums devoted to specific interests or topics [22]. On the other hand, technically Reddit is not a true social network, as it does not allow users to select others as friends and receive information from them, but only to subscribe to subreddits of interest. Moreover, Reddit users are only weakly influenced by recommendation algorithms, due to the platform's limited use of personalized recommendations. This means that all users accessing a given subreddit simultaneously are exposed to the same content, unlike what happens on Facebook or X – especially when echo chambers are present. These specific characteristics may render existing analytical approaches inadequate for detecting echo chambers on Reddit.

Here we identify meaningful patterns of echo chambers on Reddit political discourse by exploiting statistical validation techniques, which increase the signal-to-noise ratio by reducing the density and dimensionality of user interactions [23, 24]. We focus on three key aspects of users' political behavior: patterns of ideological alignment, habits of information consumption, and dynamics of political polarization. These features are uncovered using a dual-layer network approach: i) a "user-interaction" network connecting users to the subreddits they contribute to, characterizing their polarization in terms of shared user bases; ii) an "information-diet" network linking subreddits to external domains through news links shared within their posts, revealing information dissemination patterns across communities. We extract statistically reliable structures from these networks, label them with subreddit-level tags assigned from metadata, trace their evolution over time and characterize their overlap to detect echo chambers.

Our analysis focuses on political discourse during Trump's 2016 campaign and post-election period. We find that overall political polarization on Reddit declines between 2013 and 2017. However, certain communities – especially Democratic and Conservative groups (with the latter term referring to broader center-right movements beyond the Republican Party) – remain consistently polarized. Furthermore, users from banned communities, which initially aligned with far-right groups, shift toward

conservative subreddits during the 2016 elections, continuing to influence conservative discourse even after the ban of their original communities. Linguistic analysis reveals that users adopt increasingly similar language patterns within ideological groups over time. Moreover, communities sharing similar information are also those with similar user base, reinforcing echo chamber formation where users interact within closed ideological environments while consuming identical information sources. Finally, despite increased cross-group interactions during major political events, declining comment scores between opposing groups suggest that greater exposure paradoxically coincides with increased antagonism rather than reduced polarization.

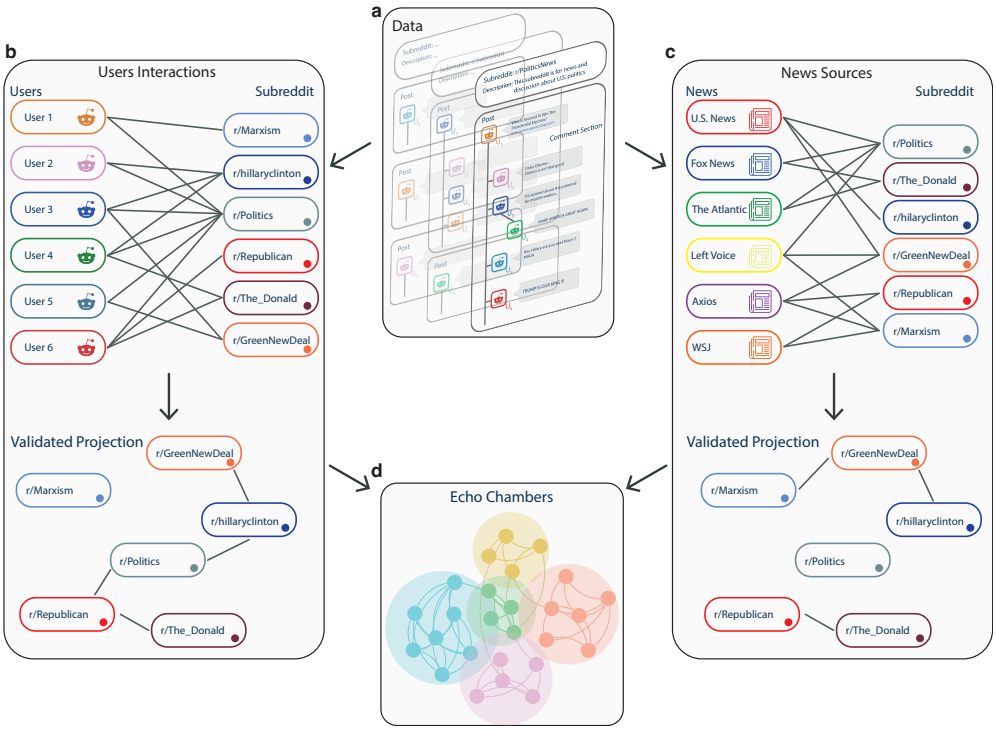

Fig. 1: **Methodological framework**. This figure summarizes the main approaches used in this study, analyzing two dimensions–news domains and user interactions–to reveal the presence of echo chambers. **a** Structure of Reddit data: conversation trees with root posts and comments underneath, belonging to different politics-related subreddits. **b** Construction of the "user-interaction" network: from the bipartite network of users interacting on subreddits to the validated network of significant similarities among subreddits (in terms of common user base). **c** "Information-diet" approach: from the bipartite network of news domain shared by subreddits to the validated network of significant similarities among subreddits (in terms of commonly shared domains). **d** Matching between the communities of subreddits in the two validated networks reveals echo chambers: closed discussion forums characterized by overlapping user and news source patterns.

Results

The structure of a conversation on Reddit consists of an initial post (often containing a question, opinion, news item, picture, video, etc.), followed by a tree-like series of users' comments to the original post or to other comments in the thread (Figure 1a). Our dataset comprises 120.4M comments made by 1.9M Reddit users on a set of 498 politically oriented subreddits from January 2013 to December 2017, obtained from the Reddit Politosphere dataset [25], as well as the corresponding 16.8M original submissions (posts) obtained from the Pushshift Reddit Dataset [26], sharing in total a set of 208.3K news domains (e.g., cnn.com, foxnews.com, theguardian.com; see Supplementary Information, **S1** for further details and general temporal trend of the data). We use this information to build two types of network, capturing the similarity among subreddits in terms of user interaction and news consumption patterns, respectively (see Methods). The first, the "user-interaction" projection, focuses on how users co-participate in communities, highlighting patterns of ideological alignment. The second, the "information-diet" projection, captures similarities in news consumption across communities.

In the "user-interaction" approach (see Figure 1b), we build a bipartite network of users and subreddits, with weighted links measuring how many times users contributed to each subreddit. We then project it on the users layer to obtain a monopartite network of subreddits, connected according to how many users participate in both their discussions. In the "information-diet" approach (Figure 1c), we build a bipartite network of news domains and subreddits, with weighted links measuring how many times domains are shared within the posts of the various subreddits. We then project it on the domains layer to obtain a monopartite network of subreddits, this time connected according to how many times they contain posts sharing the same domain. In both cases, raw co-occurrences are biased by user activity, as well as subreddit and domain popularity. To correct for these effects, we apply rigorous statistical filtering based on maximum-entropy null models [23], which retain only the most significant, non-random connections (see **Methods** and Supplementary Information, **S2**). This approach, which has been successfully applied in the past across domains [27–29], including analyses on Twitter [14, 30–32], ensures that the resulting projections retain only statistically significant co-occurrence patterns, filtering out artifacts of individual heterogeneity and random fluctuations.

After such statistical validation, we perform community detection using the Louvain algorithm [33]. Indeed, the validation process substantially sparsifies the networks, improving the signal-to-noise ratio and allowing for the identification of robust communities of similar subreddits. For instance in the user–subreddit projection, modularity increases and density decreases by roughly an order of magnitude with respect to the unfiltered network, while removing only 16% of nodes. Note that communities of the statistically validated networks are robust with respect to the choice of the community detection method [34, 35] (see Supplementary Information, **S3**), while without validation communities become unstable and highly sensitive to parameter choices (see Supplementary Information, **S2**). According to Ref. [12], echo chambers emerge when patterns of homophily, i.e. users holding similar opinions, align with patterns of polarization, measured through the similarity in the news domains linked in their

shared content. Translating this operational definition to our context here we detect echo chambers as the overlaps between communities detected with the two approaches described above, i.e. the "user-interaction" and "information-diet" ones. In practice, in the present application, echo chambers are groups of subreddits populated by the same set of users who share contents from the same set of information sources (see Figure1d).

In Politosphere, subreddits are tagged as "Democratic" or "Conservative" according to their political leaning, while "Banned" subreddits receive a dedicated tag. We extend this scheme using additional tags that we manually assign to subreddits, based on their titles, public descriptions, and moderator notes (see Supplementary Materials **S1**). These tags further enable group-level analyses, as communities are assigned tags according to the tag frequencies among their constituent subreddits. In the following plots, each tag is assigned a base color; the color of a community is then computed as the average of the colors of its constituent tags (see Supplementary Materials **S1**).

User-interaction network: Communities of subreddits with similar user bases

Reddit's political communities, based on user–interaction analysis, organize into ideologically and thematically aligned clusters, with persistent groups alongside others showing partial mixing or fragmentation over time (see Figure 2**a**). In 2013 we observe distinct clusters: a Far-Left group (light blue), a Conservative–Libertarian cluster (pink), News and Social Justice subreddits (green and orange), together with thematic communities focused on Geopolitics (khaki), Canada (yellow), and the UK/Europe (brown) (see panel 2**b** for the corresponding communities highlighted in the network). In 2014, coinciding with the U.S. midterm elections, a similar partition persists, with the first appearance of a joint Ban/Far-Right cluster (dark purple), a distinct Libertarian group (pink), and the consolidation of Social Justice and Political Talk communities. Democratic and Conservative subreddits belong with a mixed cluster. By 2015, the structure becomes richer: the Far-Left cluster remains stable, Geopolitics separates from News communities, and the Canada group persists. Democratic and Conservative subreddits still co-exist within the same cluster, though their separation becomes increasingly visible (see Figure 2 **c**).

The 2016 U.S. presidential election represents a turning point (see Figure 2 **c**). A Democratic cluster (blue) emerges more clearly, closely connected with News and Geopolitics communities. Conservative subreddits (red) form their own cluster, often in proximity to Libertarian, Gun Rights, and Economic groups (e.g., r/Elections, r/Conservative, r/KasichForPresident, r/progun, r/CatholicPolitics, r/climateskeptics). Hybrid communities also appear, merging pro-Trump Conservatives, Far-Right groups, banned subreddits (e.g., r/WhiteRights, r/fascism, r/The_Donald), and even some Social Justice spaces – reflecting the alignment of distinct ideological currents within shared spaces. In 2017, the division in Democratic and Conservative clusters further consolidate. Democrats form a stable community, while Conservatives, already split in 2016, remain divided across clusters: some align with more moderate or mixed political subreddits, while others persist within Trump-centric and hybrid communities that had emerged the year before. Ban and Far-Right groups continue to coalesce,

overlapping with both Social Justice and Conservative subreddits. Meanwhile, smaller but persistent clusters focused on Geopolitics, Far-Left politics, and non-U.S. topics (Canada, Australia, UK) maintain continuity throughout the period. These shifts are discussed in further detail in Supplementary Materials **S4** and details on community topic labeling are provided in **S1**.

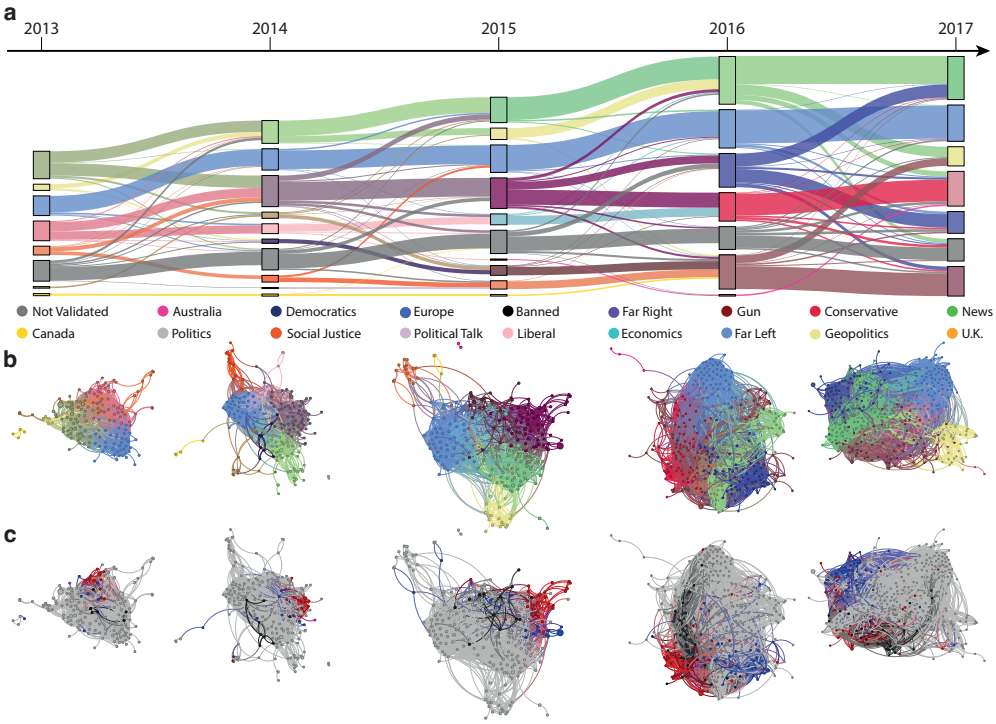

Fig. 2: **Temporal evolution of subreddits' communities with similar user base**. **a)** Sankey diagram where the width of the flows is proportional to the number of subreddits moving among communities. Community colors reflect the composite hues of their underlying subreddit tags, shown in the legend. Label "Not Validated" refers to subreddits that are not connected to others by statistically validated links. **b)** Evolution of the main communities of the validated subreddits, with several communities exhibiting strong tag homogeneity. **c)** Focus on Democratic, Conservative, and Banned subreddits. As elections approach, both Democratic and Conservative subreddits increase in number, while Banned subreddits exhibit more frequent co-occurrences with Conservative ones.

Polarization, banned users and cross interactions

Having established the validated community structure of political interactions, we analyze polarization patterns and the evolution of community composition over time. We

compute polarization from two complementary perspectives: (i) tag-based communities, i.e., groups of subreddits sharing the same thematic tag; and (ii) network-detected communities, obtained from validated projections of the user–subreddit interaction network. To study the level of polarization of the political debate, we label users with the index of the community where they are relatively more active (see Methods for further details).

Figure 3**a** shows donut charts for selected tag-based groups (Far-Left, Democratic, Conservative, and Far-Right) in 2014 and 2016, mapping the full composition of user labels within each group. Figure 3**b** displays the corresponding distributions for a set of detected network communities (composition examined in terms of user tags) that most closely align with the same categories. Additional years and full details on the construction of these charts are reported in Supplementary Information, **S4**. Both perspectives reveal that most communities are dominated by one or a few user labels, indicating that the same groups of users tend to concentrate activity within specific forums. The comparison between 2014 and 2016 shows that although overall polarization tends to decline over time, certain groups remain persistently polarized – particularly Democratic-aligned communities and banned subreddits, including conservative factions associated with Trump. These patterns point to sustained divisive dynamics among conservative-aligned groups, with banned users playing a disproportionately influential role in shaping their evolution. We also observe a temporal realignment involving Far-Right and Banned users: initially, Far-Right groups overlap mainly with Banned tags, but over time they become increasingly embedded within mixed Conservative–Banned–Social Justice communities (see Supplementary Information, **S4** for the full set of donut charts, including Banned-tag distributions).

Additional evidence of polarization is provided in Figure 3**c**, in which we analyzed selected communities in tag-based and interaction-based approaches. We consider the polarization index $\rho \in [0, 1]$ which measures the extent to which a community's user base is self-focused, with higher values indicating stronger polarization (see Methods). Polarization emerges from highly mixed communities combining Far-Right, Banned, Conservative, and Social Justice subreddits, with these heterogeneous clusters remaining strongly polarized despite their diverse composition. Banned communities display particularly sharp increases in polarization over time–rising nearly sixfold–while Conservative communities maintain high polarization levels regardless of explicit labeling.

Figure 3d reports average comment scores normalized by subreddit size for 2014 and 2016, computed separately for tag-based communities (top) and interaction-based communities (bottom). Each comment's score is first normalized by the number of active users in the corresponding subreddit to correct for visibility bias, and then aggregated within each community as the total normalized score divided by the total number of comments. As elections approach, Democratic, Conservative, and Banned communities exhibit declining scores even as cross-community interactions intensify. Groups such as Social Justice and Far-Right show higher scores when clearly distinct but experience sharp drops as ideological boundaries blur, whereas Far-Left communities remain comparatively stable, reflecting a more consistent ideological stance. Overall, these trends suggest that increasing cross-group participation coincides with

reduced internal cohesion and diminished group distinctiveness. Figure 3e shows linguistic similarity patterns based on cosine similarity of text embeddings. Subreddits sharing the same tags exhibit high internal similarity, particularly during election years, while in 2016 a convergence in language use between Democratic, Conservative, and Banned subreddits indicates increasingly blurred ideological boundaries despite persistent within-group cohesion. A more rigorous analysis, reported in Supplementary Section **S6**, confirms these results and reveals a statistically significant rise in linguistic similarity over time, especially among Conservative, Democratic, and Far-Left subreddits.

Information Ecosystem and Echo Chambers

We now take the "information-diet" approach to understand how Reddit communities consume and share news. Similarly to the "user-interaction" network, the validated projection has a marked community structure (see Figure4**a**), witnessing the existence of groups of subreddits that tend to rely on similar sets of news sources (see Supplementary Information, **S7** for network statistics). The flow diagram of Figure 4a shows how domain-sharing communities evolve year by year. For instance, Democratic- and Conservative-aligned communities emerge clearly, particularly after the 2016 election, when they begin to rely on increasingly distinct domains. Communities also form around geographic interests, including Canada, Australia, and UK. A group of "Banned" subreddits, initially close to far-right spaces, moves closer to Conservative communities in the years following the election. On the left side, different subgroups emerge with specific focuses: one on social justice issues, another on Marxism and Anarchism, and a third centered on meme-sharing. Additional analyses on domain-level polarization and label propagation across domain categories are provided in Supplementary Section **S7**, with full details on the networks and community structures for both approaches in **S8**. Notably, several of these groups match those observed in the "user-interaction" networks (see Figure 4**b**).

To further explore the alignment between "user-interaction" and "information-diet" communities at the user level, we construct a bipartite network where each node represents a community from one of the two approaches, while links are weighted by the number of active users (see Methods). We can then assess statistically significant overlaps using the BiWCM null model (see Methods). Although community matches emerge in correspondence with similar topics, even with very low p-values, few entries remain statistically significant once correcting for multiple hypothesis testing with the false discovery rate (see heatmaps of community matches in Supplementary Materials **S9**).

A more fine-grained analysis on sub-communities – obtained by using the Louvain algorithm again on individual communities – reveals clearer patterns. Communities that focus on the same topics (see Supplementary Materials **S1**), such as gun rights or economic policy, are much more likely to share both users and information sources, as confirmed by the statistically significant overlaps (see Tab. 1 for p-values). These correspondences suggest that ideological coherence is particularly strong in more focused environments.

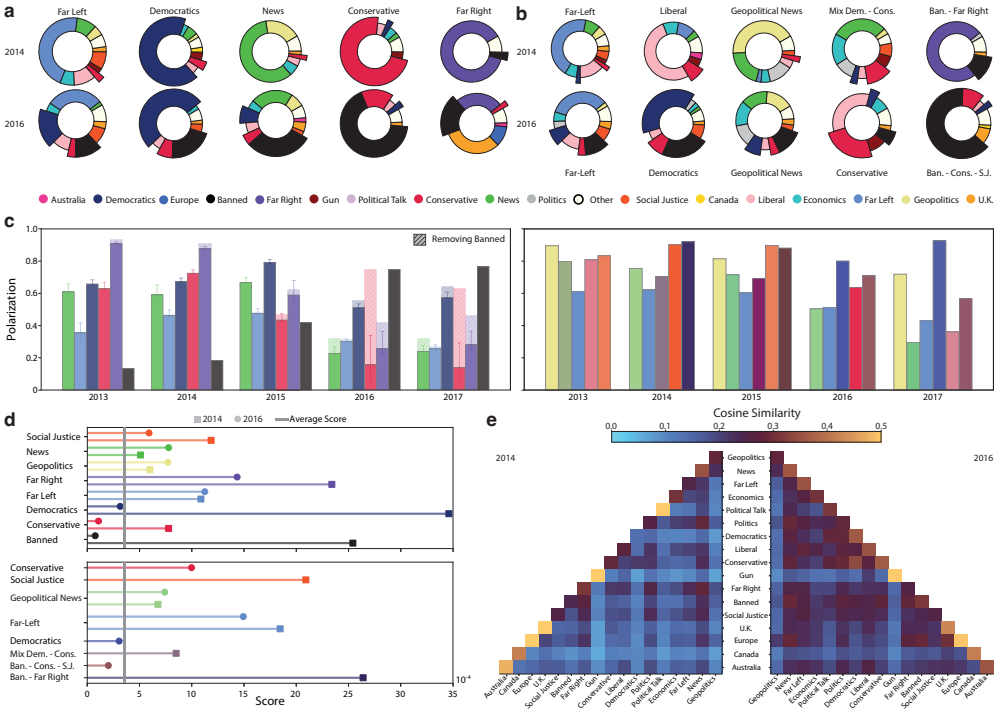

Fig. 3: **Polarization, engagement and linguistic similarity across subreddit communities**. **a** Donuts charts showing the distribution of user labels within the communities of the validated "user-interaction" network projections, for 2014 and 2016. **b** Donuts charts for tag-based communities, for 2014 and 2016. In both cases, polarization decreases in time, with some exceptions (e.g., Democratics and Banned). **c** Bar plots showing annual polarization levels across a selection of communities, with solid bars computed for all users and dashed bars computed excluding "Banned" users. This distinction allows to capture the strong polarizing role of banned subreddits, especially within conservative clusters. Early polarization is most pronounced among Far-Right and Democratic groups, with major shifts observed in the Conservative category. Over time, overall polarization decreases, except for a notable rise among banned communities. **d** Average scores of comments for 2014 and 2016, revealing a general decline in user engagement, particularly in Democratic, Conservative, and Banned subreddits. **e** Cosine similarity between posts within topic-based subreddit groups, indicating strong intra-group similarity and increasing inter-group similarity over time, especially between Democratic and Conservative communities.

These patterns challenge the notion that echo chambers, on Reddit or elsewhere, represent only a marginal phenomenon involving a small fraction of users [14, 19, 20]. Here, we find many groups of subreddits that significantly share both user bases and news sources. Between 40% and 60% of users belong to such overlapping communities, while the proportion of users active within validated sub-community matches increases

from about 20% in 2013–2014 to 36% in 2015, remaining stable in 2016 (28%) and 2017 (34%). This pattern highlights a marked consolidation of echo-chamber structures, particularly during politically salient years such as 2014 and 2016. Together, these findings highlight how user engagement and information sharing reinforce each other in shaping Reddit's political landscape. Full details on community structure, overlap sub-communities patterns, and temporal dynamics are available in the Supplementary Information, **S8-S9**.

Focus on Democratic, Conservative, and Banned Groups

We finally perform a closer investigation of the political discourse between Democratic and Conservative communities, as these factions emerge distinctly across the network approaches we considered. Figure 5a shows yearly distances between the linguistic patterns of these groups, computed from subreddit-level text embeddings (see Methods). We observe a decrease in the distance (from 0.8 to 0.7) between pairs of Democrats, Conservatives, and banned users during the election cycle. As benchmarks, we consider a random model (mean similarity of randomly assembled communities of the same size) and a heterogeneous model (mean similarity of all pairwise subreddits within a community), which confirm that the observed convergence exceeds what is expected from random mixing or compositional effects (see Methods). The convergence between Democrats and Conservatives is particularly evident in 2016, likely driven by the shared focus on election-related narratives (see Figure 5**a**).

We then consider the average distance between these communities (see Methods) computed on the statistically validated "user–interaction" network. Figure 5b shows a marked increase in the average network distance between Democratic and Conservative subreddits starting in 2015 (the year preceding Trump's election) and peaking in 2016. This separation is particularly pronounced for Democratic subreddits, which grow increasingly distant from their Conservative counterparts. Statistical tests confirm the significance of this trend when comparing Democrat–Conservative pairs with previous years and with overall network distances. A similar pattern is reflected in the average distance computed on the "information-diet" validated networks, shown in Figure 5c: Democratic and Conservative communities initially share many domains, but their media ecosystems gradually diverge in the lead-up to the election. Conservatives remain consistently closer to Banned subreddits than Democrats throughout the period.

To better understand cross-group dynamics, we track user-level interactions by filtering the original dataset of comments, assigning labels to users through propagation from subreddit to user activity, and then retaining only comments authored by users tagged as Democratic, Conservative, or Banned, regardless of the subreddit. Figure 5d shows the proportion of comments exchanged within and between these groups (here reported for 2014 and 2016), measured as the normalized number of comments exchanged between groups relative to the total number of comments produced by users of each group. Over the period, the total number of comments by these groups grows substantially, reaching nearly 10^6 in 2016, about ten times more than in 2014–while cross-group commenting steadily increases but is accompanied

by a marked decline in comment scores. Average scores drop by 93.4 % in Democratic subreddits and by 85.6 % in Conservative ones, with the decline most severe in candidate-centered communities such as r/SandersForPresident, r/HillaryClinton, and r/The_Donald, while remaining comparatively stable in less polarizing spaces (e.g., r/KasichForPresident). As shown in Figure 5e, cross-group interactions are evaluated even more negatively than within-group ones. Notably, the number of Democratic and Conservative users engaging in each other's communities rises over time, independently of overall volume, suggesting growing exposure to opposing viewpoints even as their reception worsens (a trend consistent with prior literature [17, 19, 20]). Full results across all years, including this breakdown, and further analyses are provided in Supplementary Information **S10**.

Discussions

Online social networks data often come with a high level of noise, due to the multifacet nature of social interactions, their large heterogeneity and the intrinsic randomness of human behavior [24]. To increase the signal-to-noise ratio, we applied a statistical filtering approach to isolate the core structure of political discourse, enabling reliable detection of persistent patterns of political polarization and echo chambers on Reddit [16, 36]. We focused on the US political debate from 2013 to 2017, uncovering a robust architecture of ideological communities that endures over time. Despite widespread interaction across groups, polarization remains deeply embedded – offering novel perspectives into how echo chambers sustain themselves even amid moments of heightened engagement and debate.

Within the observed clusters, users are repeatedly exposed to a narrow set of information sources and engage with the same topics and individuals, often using highly similar language, suggesting that echo chambers are not limited to fringe groups and are present even during periods of apparent openness. Remarkably, partisan communities showed a decline in average comment scores, especially in candidate-focused forums during election periods, a trend also observed in r/The_Donald [37], suggesting increasing internal heterogeneity even within apparently cohesive groups. Finally, during election periods, linguistic convergence across Democrat, mainstream Conservative, and even Banned communities suggested intensified cross-camp debate without genuine integration, as increasingly similar language hinted at shared rhetorical frames across ideological boundaries [38]. Our findings are mirrored in real-world political tensions, as the clustering of communities and semantic distances across topic groups reveal systematic contrasts between Democratic and Conservative spheres. Remarkably, both factions remain equally detached from economics- and social-justice-related topics – echoing post-election critiques [39] and concerns over the marginalization of economically "left-behind" U.S. communities [40].

The observation of a consolidated Conservative cluster, heavily influenced by users of banned communities, resonates with prior findings that the 2016 polarization spike was distinctly asymmetric and driven by right-wing activity [21]. However, Ref. [21] employs neural community embeddings [41], defining community similarity in terms

of user bases similarity within a high-dimensional space. Here instead we define validated network communities in terms of user interaction (whatever their traits) and information sources – a crucial dimension to define echo chambers. These different constructions can explain why, while individual users rarely alter their personal ideological trajectories over time [21], their structural integration can profoundly alter the social environment.

We tested our procedures on several dimensions to detect possible weaknesses. First, our approach to user classification, based primarily on participation metrics, rather than direct textual ideological content, may introduce some misalignments. Such discrepancies can arise when participation patterns evolve rapidly, thereby reducing their correspondence with users' underlying political orientations. Tracking these users across subsequent years revealed frequent shifts toward Conservative or Banned communities, followed in some cases by partial returns to Democratic spaces – a pattern likely driven by adversarial interactions rather than genuine ideological shifts (see Supplementary Sections **S2, S10**). Second, differences in temporal resolution may influence the granularity of observed dynamics, since zooming into shorter periods can reveal transient variations in community composition and polarization intensity. We verified the robustness of our results using four-month aggregation windows instead of yearly ones, obtaining consistent community structures and polarization trends, with minor differences (see Supplementary Information **S11**). Third, differences in level of aggregation may also account for discrepancies across studies. For instance, while our analysis operates at the level of domains and subreddits – thus, focusing on aggregated community behavior – Ref. [14] examines news consumption and polarization directly at the level of individual users and URLs. This user-centric perspective could capture interaction patterns at a different scale, thereby also leading to variations in the measured strength of echo-chamber effects across systems.

Future developments of our research should enhance classification accuracy by integrating sentiment and textual analyses of user comments, moving beyond participation metrics alone — an increasingly relevant step given the growing presence and human-like behavior of automated and AI-driven accounts in online discussions [42–44]. Investigating the mechanisms behind linguistic similarity during election cycles – particularly their links with linguistic complexity and user involvement within echo chambers [45] – represents a promising direction for understanding the evolution of political discourse. Furthermore, dedicated analyses of highly polarized communities, such as those explicitly linked to Trump, may uncover distinctive interaction patterns. Additional work is needed to clarify the influence of Banned users within Conservative-aligned clusters and to assess the effectiveness of moderation strategies – particularly subreddit bans – in mitigating long-term polarization [46, 47].

Understanding echo chambers as persistent, self-sustaining structures – rather than transient anomalies – shifts the focus from merely facilitating cross-group exposure to designing interactions that promote reflection over reaction. Whether platforms will embrace such responsibility remains an open question, but our results indicate that current dynamics are unlikely to change spontaneously.

Methods

Dataset Description

We used the Reddit Politosphere dataset [25], which collects all comments from a large set of politically oriented subreddits between 2013 and 2017, and complemented it with the corresponding submissions from the Pushshift Reddit dataset [26]. Together, these data comprise millions of comments and posts produced by hundreds of subreddits and domains over the five-year period (see Supplementary Information, Section S1, for details and visualizations).

Each record comes with rich metadata. For comments, available fields include user and submission identifiers, the subreddit in which the comment appears, parent comment ID, timestamps, scores, and the comment text. In addition, the dataset creators provide curated metadata, such as information about the subreddit (e.g., whether it was later banned, or associated with Democratic or Conservative communities), along with other annotations. For submissions, metadata include author and submission IDs, timestamps, scores, textual content, and, when applicable, the external URL and the domain of the linked content. Subreddit-level metadata about political alignment were further extended into a set of topic tags covering 17 categories, allowing for broader classification across all subreddits. The manual tagging procedure was validated using GPT-4. We also used the public description of each subreddit to assist in the manual labeling of subreddits into the corresponding topic tags (see Supplementary Information, **S12**).

Network validation

The validation framework starts from two weighted bipartite networks representing Reddit interactions: (i) users commenting on subreddits, and (ii) domains shared within subreddits. Here we discuss the case of the users-subreddits network, but the discussion applies to the domains-subreddit network as well.

We extract statistically meaningful projections onto the subreddit layer by identifying significant co-occurrence patterns [27, 48]. Prior to validation, we binarize networks using the Revealed Comparative Advantage (RCA) [49], which compares the observed weight of each link to the expected activity of the two nodes involved, thus correcting for heterogeneity in activity levels. Only links with $RCA \geq 1$ are retained in the binarized network. After binarization, we project onto the subreddit layer by counting common users between subreddits. In particular, if $b_{i\alpha}$ is the generic entry of the (binarized) biadjacency matrix between subreddit i and user α, the expected number of common users between two subreddits i and j is $V_{ij} = \sum_{\alpha} b_{i\alpha} b_{j\alpha}$. However, this raw overlap is biased by subreddit popularity and user activity, requiring further statistical validation.

To correct for such biases, we apply statistical validation based on maximum entropy null models [23]. In particular, we use the Bipartite Configuration Model (BiCM) [50] based on constraining degrees of both classes of nodes (subreddits and users). According to BiCM, co-occurrences V_{ij} are modeled with a Poisson–Binomial

distribution, i.e., the sum of independent Bernoulli trials with different success probabilities - which depend on the degrees of the the two subreddits and of the individual users involved. The corresponding p-value is

$$p[V_{ij}] = 1 - \sum_{x=0}^{V_{ij}-1} \pi(x \mid i, j), \tag{0.1}$$

where $\pi(x \mid i, j)$ denotes the probability of observing exactly x shared neighbors according to the BiCM null model.

We validate co-occurrences by applying the False Discovery Rate (FDR) correction [51] to the set of computed p-values. Let D be the number of hypotheses tested and α the desired significance level. Ordering the p-values in increasing order, we select the largest one satisfying $p\text{-value}_i \leq i\alpha/D$, where i is the rank of the p-value. All links with p-values below this threshold are retained in the validated projection, filtering out spurious co-occurrences and isolating statistically robust similarities between subreddits [27].

Polarization index

For a given partition of subreddits into communities, we label users according to where they are relatively more active, correcting for heterogeneity in community size. Specifically, for each user i and community c, let n_{ic} be the number of subreddits in c where i has commented and S_c the total number of subreddits in c. The user label is $u_i = \arg\max_c(n_{ic}/S_c)$ [30].

Let U_c be the set of users labeled as c, and W_c the set of users who commented at least once on subreddits in community c. The polarization index of community c is $\rho_c = |W_c \cap U_c|/|W_c| \in [0, 1]$, i.e., the fraction of active users in community c whose main activity is concentrated in c. Higher values indicate stronger polarization.

In addition to this index, we also analyze the full composition of communities, i.e., how users with different labels populate each community. This information is visualized in the donut charts of the main text and corresponds to the full polarization matrices described in Supplementary Information **S4**.

Text Preprocessing and Similarity Analysis

To compute textual similarity, the texts of Reddit posts are first preprocessed by converting all content to lowercase, removing stopwords, and applying lemmatization. From this cleaned corpus, we constructed a vocabulary of unigrams, bigrams, and trigrams. To ensure comparability across subreddits, we retained only those n-grams that appeared more than once in the overall corpus.

To represent textual content numerically, we used the standard Term Frequency–Inverse Document Frequency (TF-IDF) representation [52]. This approach weighs terms by how often they appear within a subreddit while reducing the weight of terms that are frequent across many subreddits, thereby highlighting words that are more distinctive and informative for each community.

Following vectorization, textual similarity between documents was computed using cosine similarity:

$$\text{CosineSim}(A, B) = \frac{A \cdot B}{|A| \cdot |B|} \tag{0.2}$$

where A and B are the TF-IDF vectors of two documents. This metric, commonly used in information retrieval [52], captures how aligned two documents are in the vector space, independently of length.

Note that in Fig. 5 we report cosine distances, defined as $1-\text{CosineSim}(A, B)$. This choice was made to obtain a true distance metric, allowing for direct comparison with network-based distances (e.g., path lengths), thus making the cross-modal analysis more consistent.

Analysis and validation of Echo Chambers

To analyze echo chambers, we compare subreddit communities derived from "user-interaction" networks and those obtained from "information-diet" patterns. Specifically, we construct overlap matrices where rows correspond to interaction-based communities and columns to domain-based communities. Each matrix entry reflects the number of subreddits shared between the corresponding pair of communities. To assess the statistical significance of these overlaps, we constructed a bipartite network in which nodes on each layer represent the communities from the two classifications. A link connects two nodes if the corresponding matrix entry is nonzero, and its weight is given by the number of users active in the set of shared subreddits, capturing user engagement in the shared subreddits, a suitable parameter for studying overlaps between groups in the context of echo chambers.

This bipartite structure was validated using the Bipartite Weighted Configuration Model (BiWCM) [53], a maximum entropy null model for bipartite weighted networks that constrain values of node strengths. To control for multiple hypothesis testing, we applied the False Discovery Rate (FDR) correction [51] introduced above to the p-values obtained from the BiWCM validation. To increase the resolution of the analysis and address the limitations imposed by the relatively small number of original communities, we repeated the validation procedure using subcommunities. Specifically, for each community in the validated subreddit networks, we extract the corresponding subgraph and identified finer-grained substructures using again the Louvain algorithm. These subcommunities were then used to construct new bipartite overlap matrices, subsequently validated using the BiWCM model.

The entries found to be statistically significant in these higher-resolution matrices are further tested using the two-sided Mann–Whitney U test, comparing the distribution of values in validated entries against all other matrix entries: the test confirms that validated overlaps follow a significantly different distribution, supporting the robustness of the results.

Network distance between communities

To define a distance between two communities in the validated subreddit networks, we compute the average pairwise distance between all subreddit pairs belonging to

the two communities and present in the statistically validated projection. Given the possibility that some pairs of subreddits may belong to disconnected components of the validated network, we define the distance between two nodes (subreddits) i and j as the reciprocal of the shortest path length $\mathrm{sp}(i,j)$ between them: $d_{ij} = 1/\mathrm{sp}(i,j)$ if a path exists between them, and $d_{ij} = 0$ otherwise. The average distance between two communities c and c', defined by sets of subreddits S_c and $S_{c'}$, is then computed as the harmonic mean of all pairwise distances:

$$D_{CC'} = \frac{|S_C| \cdot |S_{C'}|}{\sum_{i \in S_C} \sum_{j \in S_{C'}} d_{ij}^{-1}} \tag{0.3}$$

To account for variations in network structure over years, each distance $D_{CC'}$ is normalized by the average $\langle D \rangle$ of all d_{ij} values (for all pairs of connected nodes) in the validated network for that year: $\hat{D}_{CC'} = D_{CC'}/\langle D \rangle$. Finally to assess the statistical significance of changes in distance between communities, we use a two-sided Kolmogorov–Smirnov (KS) test. Specifically, for each year, we compare the distribution of pairwise distances (d_{ij}) between two communities with the same distribution in the previous year, and the full distribution of all pairwise distances in the current network.

Declarations

- Data Availability: Reddit conversation data used in this study include comments from the Politosphere dataset [25], originally collected via the Pushshift API (https://www.reddit.com/r/pushshift/), and posts retrieved from the same API in 2022.
- Code Availability: The code used to reproduce the analysis is available on GitHub. For inquiries or additional material, please contact D. Cirulli at daniele.cirulli@cref.it.
- Acknowledgements: A.D. and D.C. acknowledge the helpful discussions on Reddit analysis with Anna Mancini and Riccardo Di Clemente.
- Funding Declaration: G.C. acknowledges support from "Deep 'N Rec" Progetto di Ricerca di Ateneo of University of Rome Tor Vergata. F.S. was partially supported by the project "CODE – Coupling Opinion Dynamics with Epidemics", funded under PNRR Mission 4 "Education and Research" - Component C2 - Investment 1.1 - Next Generation EU "Fund for National Research Program and Projects of Significant National Interest" PRIN 2022 PNRR, grant code P2022AKRZ9. The work of F.S. is part of the joint initiative CREF-SONY.
- Author Contributions: D.C. and A.D. gathered the data. D.C. performed the analysis. D.C. and A.D. made the figures. F.S. and G.C designed the analysis and supervised the project. All authors discussed the results and contributed to the final manuscript.
- Competing Interests: The authors declare no competing interests.
- This work makes extensive use of open-source libraries and tools. For data handling and analysis, we used Pandas[54], NumPy[55], SciPy[56], NLTK[57] and Scikit-learn[52]. For network analysis and visualization, we relied on NetworkX [58],

Interaction Community	Information Community	p-value
Geop/News-	Dem/Cons-	2e-08
Far-Left	Far-Left	7e-04
Far-Left	Lib/Dem-	7e-05
Far-Left	Far-Left	3e-04
Far-Left	Econ/Far-Left-	8e-04
PolTalk	Austr/Far-Left-	5e-06
Dem/SocialJustice-	Far-Left/SocialJustice	5e-06
Dem/Politic	Dem/Cons-	3e-03
Dem/Politic	Lib/Cons-	3e-05
Dem/Cons-	Far-Left	4e-05
Dem/Cons-	Far-Left/Politic	1e-06
Lib	Austr/Far-Left-	2e-04
Cons	Far-Left/Geop-	3e-07
Gun	Cons/Gun-	3e-03
Far-Right/Econ-	Cons/Gun-	1e-07
Ban	Ban/Geop-	3e-13
SocialJustice	News/Politic	1e-06
SocialJustice	SocialJustice/Politic	6e-08
SocialJustice	Econ/Far-Left-	1e-03
SocialJustice	SocialJustice/Politic	6e-05
UK	UK	2e-09
Eur	News/Politic	2e-04
Eur	Austr/Far-Left-	1e-04
Canada	Canada	1e-11
Austr	Austr/Far-Left-	3e-09

(a) 2014

Interaction Community	Information Community	p-value
Geop	Geop/News	2e-05
Geop	Geop	1e-08
News	News/Politic	7e-06
News/Politic	Geop/News	8e-06
News/Econ-	News/Politic	2e-05
News/Econ-	Econ/Far-Left	5e-04
Far-Left	Far-Left	9e-07
Far-Left/SocialJustice-	Far-Left	2e-06
Far-Left/SocialJustice-	PolTalk	5e-04
Far-Left/SocialJustice-	SocialJustice	5e-06
Far-Left	Far-Left	3e-08
Far-Left	Econ/Far-Left	2e-04
Far-Left	Far-Right/Far-Left-	2e-04
Dem/News-	Far-Left/News-	4e-05
Dem/News-	Dem/Cons-	5e-04
Lib/Econ-	Far-Left/Politic	5e-04
Cons	Cons/Politic	1e-04
Cons/Ban	Dem/Ban-	9e-05
Gun	Gun/Dem-	2e-07
Ban/Far-Right-	Ban/Far-Right	7e-09
SocialJustice/PolTalk-	Far-Left	1e-05
SocialJustice/PolTalk-	PolTalk	4e-05
SocialJustice/PolTalk-	Ban/Far-Right	4e-04
UK/Eur	Canada	6e-09
Canada	Canada	2e-15
Austr	Far-Left/Austr-	4e-12

(b) 2016

Table 1: Significance of echo chamber correspondences in 2014 and 2016. Each list reports p-values for matches between subreddit subcommunities detected from user–subreddit interactions and from subreddit–domain sharing. All values shown correspond to statistically significant overlaps after False Discovery Rate correction ($\alpha = 0.05$).

iGraph[59], and graph-tool[60], while Gephi[61] was employed for network visualization and manual exploration. Plots and visual representations were generated using Matplotlib[62], D3Blocks[63] for chord diagrams, and Plotly[64] for Sankey diagrams. Finally, we employed NEMtropy [53] to implement Maximum Entropy Null Models for the statistical validation of Reddit's complex networks, and the weighted network model from [65] for validating echo chambers.

- Tag consistency was validated by comparing manual annotations with labels generated by GPT-4 through the OpenAI API [66] (see Supplementary Section **S12**).

References

[1] Habermas, J.: The Structural Transformation of the Public Sphere: An Inquiry Into a Category of Bourgeois Society. Studies in contemporary German social thought. Polity Press, ??? (1989). https://books.google.it/books?id=3iIVnwEACAAJ

[2] Sylvester, D.E., McGlynn, A.J.: The digital divide, political participation, and place. Social Science Computer Review **28**(1), 64–74 (2009) https://doi.org/10.1177/0894439309335148

[3] Bail, C.: Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing. Princeton University Press, ??? (2021). https://doi.org/10.2307/j.ctv18zhdhg . http://dx.doi.org/10.2307/j.ctv18zhdhg

[4] Vargo, C.J., Guo, L., McCombs, M., Shaw, D.L.: Network issue agendas on twitter during the 2012 u.s. presidential election: Network issue agendas on twitter. Journal of Communication **64**(2), 296–316 (2014) https://doi.org/10.1111/jcom.12089

[5] Han, B.-C.: Infokratie: Digitalisierung und Krise der Demokratie. Matthes & Seitz, Berlin, Deutschland (2021)

[6] Pentland, A.: Social Physics: How Good Ideas Spread—The Lessons from a New Science. The Penguin Press, New York, NY (2014)

[7] Garrett, R.K.: Echo chambers online?: Politically motivated selective exposure among internet news users. Journal of Computer-Mediated Communication **14**, 265–285 (2009) https://doi.org/10.1111/J.1083-6101.2009.01440.X

[8] Barberá, P., Jost, J.T., Nagler, J., Tucker, J.A., Bonneau, R.: Tweeting from left to right: Is online political communication more than an echo chamber? Psychological Science **26**(10), 1531–1542 (2015) https://doi.org/10.1177/0956797615594620

[9] Del Vicario, M., Vivaldo, G., Bessi, A., Zollo, F., Scala, A., Caldarelli, G., Quattrociocchi, W.: Echo chambers: Emotional contagion and group polarization on

facebook. Scientific Reports (2016) https://doi.org/10.1038/srep37825

[10] Zollo, F., Bessi, A., Del Vicario, M., Scala, A., Caldarelli, G., L. Shekhtman: Debunking in a world of tribes. PLoS ONE **7:e0181821**(12) (2017)

[11] Cossard, A., Morales, G., Kalimeri, K., Mejova, Y., Paolotti, D., Starnini, M.: Falling into the echo chamber: The italian vaccination debate on twitter. Proceedings of the International AAAI Conference on Web and Social Media **14**, 130–140 (2020) https://doi.org/10.1609/icwsm.v14i1.7285

[12] Cinelli, M., Morales, G.D.F., Galeazzi, A., Quattrociocchi, W., Starnini, M.: The echo chamber effect on social media. Proceedings of the National Academy of Sciences **118**(9), 2023301118 (2021) https://doi.org/10.1073/pnas.2023301118 https://www.pnas.org/doi/pdf/10.1073/pnas.2023301118

[13] Conover, M., Ratkiewicz, J., Francisco, M., Goncalves, B., Menczer, F., Flammini, A.: Political polarization on twitter. Proceedings of the International AAAI Conference on Web and Social Media **5**(1), 89–96 (2021) https://doi.org/10.1609/icwsm.v5i1.14126

[14] Pratelli, M., Saracco, F., Petrocchi, M.: Entropy-based detection of twitter echo chambers. PNAS Nexus **3**(5) (2024) https://doi.org/10.1093/pnasnexus/pgae177

[15] Budak, C., Nyhan, B., Rothschild, D.M., Thorson, E., Watts, D.J.: Misunderstanding the harms of online misinformation. Nature **630**(8015), 45–53 (2024) https://doi.org/10.1038/s41586-024-07417-w

[16] Falkenberg, M., Cinelli, M., Galeazzi, A., Bail, C.A., Benito, R.M., Bruns, A., Gruzd, A., Lazer, D., Lee, J.K., McCoy, J., Nagayoshi, K., Rand, D.G., Scala, A., Siegel, A., Linden, S., Varol, O., Weber, I., Wojcieszak, M., Zollo, F., Baronchelli, A., Quattrociocchi, W.: Towards global equity in political polarization research. arXiv (2025). https://doi.org/10.48550/ARXIV.2504.11090 . https://arxiv.org/abs/2504.11090

[17] De Francisci Morales, G., Monti, C., Starnini, M.: No echo in the chambers of political interactions on reddit. Scientific Reports **11**(1), 2818 (2021) https://doi.org/10.1038/s41598-021-81531-x

[18] Monti, C., D'Ignazi, J., Starnini, M., De Francisci Morales, G.: Evidence of demographic rather than ideological segregation in news discussion on reddit. In: Proceedings of the ACM Web Conference 2023. WWW '23, pp. 2777–2786. ACM, ??? (2023). https://doi.org/10.1145/3543507.3583468 . http://dx.doi.org/10.1145/3543507.3583468

[19] Morini, V., Pollacci, L., Rossetti, G.: Toward a standard approach for echo chamber detection: Reddit case study. Applied Sciences **11**(12) (2021) https://doi.org/10.3390/app11125390

[20] Colacrai, E., Cinus, F., Morales, G.D.F., Starnini, M.: Navigating multidimensional ideologies with reddit's political compass: Economic conflict and social affinity. ArXiv **abs/2401.13656** (2024)

[21] Waller, I., Anderson, A.: Quantifying social organization and political polarization in online platforms. Nature **600**(7888), 264–268 (2021) https://doi.org/10.1038/s41586-021-04167-x

[22] Reddit Inc.: Reddit Content Policy. https://www.redditinc.com/policies/content-policy. Accessed: 2024-09-24 (2024)

[23] Cimini, G., Squartini, T., Saracco, F., Garlaschelli, D., Gabrielli, A., Caldarelli, G.: The statistical physics of real-world networks. Nature Reviews Physics **1**(1), 58–71 (2019) https://doi.org/10.1038/s42254-018-0002-6

[24] De Clerck, B., Rocha, L.E., Van Utterbeeck, F.: Maximum entropy networks for large scale social network node analysis. Applied Network Science **7**(1), 68 (2022)

[25] Hofmann, V., Schütze, H., Pierrehumbert, J.B.: The reddit politosphere: A large-scale text and network resource of online political discourse. Proceedings of the International AAAI Conference on Web and Social Media **16**(1), 1259–1267 (2022) https://doi.org/10.1609/icwsm.v16i1.19377

[26] Baumgartner, J., Zannettou, S., Keegan, B., Squire, M., Blackburn, J.: The Pushshift Reddit Dataset. https://arxiv.org/abs/2001.08435 (2020)

[27] Saracco, F., Straka, M.J., Clemente, R.D., Gabrielli, A., Caldarelli, G., Squartini, T.: Inferring monopartite projections of bipartite networks: an entropy-based approach. New Journal of Physics **19**(5), 053022 (2017) https://doi.org/10.1088/1367-2630/aa6b38

[28] Cimini, G., Carra, A., Didomenicantonio, L., Zaccaria, A.: Meta-validation of bipartite network projections. Communications Physics **5**(1), 76 (2022) https://doi.org/10.1038/s42005-022-00856-9

[29] Bruno, M., Mazzilli, D., Patelli, A., Squartini, T., Saracco, F.: Inferring comparative advantage via entropy maximization. Journal of Physics: Complexity **4**(4), 045011 (2023) https://doi.org/10.1088/2632-072x/ad1411

[30] Becatti, C., Caldarelli, G., Lambiotte, R., Saracco, F.: Extracting significant signal of news consumption from social networks: the case of twitter in italian political elections. Palgrave Communications **5**(1), 91 (2019) https://doi.org/10.1057/s41599-019-0300-3

[31] Caldarelli, G., Nicola, R.D., Vigna, F.D., Petrocchi, M., Saracco, F.: The role of bot squads in the political propaganda on twitter. Communications Physics **3**, 1–15 (2020) https://doi.org/10.1038/s42005-020-0340-4

[32] Mattei, M., Pratelli, M., Caldarelli, G., Petrocchi, M., Saracco, F.: Bow-tie structures of twitter discursive communities. Scientific Reports **12**(1), 12944 (2022) https://doi.org/10.1038/s41598-022-16603-7

[33] Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. Journal of Statistical Mechanics: Theory and Experiment **2008**(10), 10008 (2008) https://doi.org/10.1088/1742-5468/2008/10/P10008

[34] Rosvall, M., Bergstrom, C.T.: Maps of random walks on complex networks reveal community structure. Proceedings of the National Academy of Sciences **105**(4), 1118–1123 (2008)

[35] Peixoto, T.P.: Hierarchical block structures and high-resolution model selection in large networks. Physical Review X **4**(1), 011047 (2014)

[36] Falkenberg, M., Zollo, F., Quattrociocchi, W., Pfeffer, J., Baronchelli, A.: Patterns of partisan toxicity and engagement reveal the common structure of online political communication across countries. Nature Communications **15**(1) (2024) https://doi.org/10.1038/s41467-024-53868-0

[37] Subtirelu, N.: Donald trump supporters and the denial of racism: An analysis of online discourse in a pro-trump community. In: Simon, A., Kerkhoff, S. (eds.) Language Aggression in Public Debates on Immigration, pp. 151–174. John Benjamins, ??? (2019). https://doi.org/10.1075/bct.102.08sub

[38] Desiderio, A., Mancini, A., Cimini, G., Di Clemente, R.: Highly engaging events reveal semantic and temporal compression in online community discourse. PNAS Nexus **4**(3), 056 (2025) https://doi.org/10.1093/pnasnexus/pgaf056

[39] Newsroom, W.: Bernie Sanders Tells Berklee Students to Create Political Art. Published November 21, 2016. https://www.wbur.org/news/2016/11/21/bernie-sanders-berklee Accessed 2025-05-03

[40] Krugman, P.: The Economics of Left-Behind Regions (2024). https://paulkrugman.substack.com/p/the-economics-of-left-behind-regions Accessed 2025-05-03

[41] Martin, T.: community2vec: Vector representations of online communities encode semantic relationships. In: Hovy, D., Volkova, S., Bamman, D., Jurgens, D., O'Connor, B., Tsur, O., Doğruöz, A.S. (eds.) Proceedings of the Second Workshop on NLP and Computational Social Science, pp. 27–31. Association for Computational Linguistics, Vancouver, Canada (2017). https://doi.org/10.18653/v1/W17-2904 . https://aclanthology.org/W17-2904/

[42] Robertson, R.E., Green, J., Ruck, D.J., Ognyanova, K., Wilson, C., Lazer, D.: Users choose to engage with more partisan news than they are exposed to

on google search. Nature **618**(7964), 342–348 (2023) https://doi.org/10.1038/s41586-023-06078-5

[43] Augenstein, I., Baldwin, T., Cha, M., Chakraborty, T., Ciampaglia, G.L., Corney, D., DiResta, R., Ferrara, E., Hale, S., Halevy, A., Hovy, E., Ji, H., Menczer, F., Miguez, R., Nakov, P., Scheufele, D., Sharma, S., Zagni, G.: Factuality challenges in the era of large language models and opportunities for fact-checking. Nature Machine Intelligence **6**(8), 852–863 (2024) https://doi.org/10.1038/s42256-024-00881-z

[44] Cirulli, D., Cimini, G., Palermo, G.: How Large Language Models play humans in online conversations: a simulated study of the 2016 US politics on Reddit. Preprint (2025)

[45] Amadori, E., Cirulli, D., Martino, E.D., Nudo, J., Sahakyan, M., Sangiorgio, E., Santoro, A., Zollo, S., Galeazzi, A., Marco, N.D.: Involvement drives complexity of language in online debates (2025). https://arxiv.org/abs/2506.22098

[46] Bond, R.M., Garrett, R.K.: Engagement with fact-checked posts on reddit. PNAS Nexus **2**(3) (2023) https://doi.org/10.1093/pnasnexus/pgad018

[47] Chandrasekharan, E., Pavalanathan, U., Srinivasan, A., Glynn, A., Eisenstein, J., Gilbert, E.: You can't stay here: The efficacy of reddit's 2015 ban examined through hate speech. Proceedings of the ACM on Human-Computer Interaction **1**(CSCW), 1–22 (2017) https://doi.org/10.1145/3134666

[48] Gualdi, S., Cimini, G., Primicerio, K., Clemente, R.D., Challet, D.: Statistically validated network of portfolio overlaps and systemic risk. Scientific Reports 2016 6:1 **6**, 1–14 (2016) https://doi.org/10.1038/srep39467

[49] Balassa, B.: Trade liberalization and revealed comparative advantage. The Manchester School of Economic and Social Studies **33**, 99–123 (1965) https://doi.org/10.1111/j.1467-9957.1965.tb00050.x

[50] Saracco, F., Clemente, R.D., Gabrielli, A., Squartini, T.: Randomizing bipartite networks: the case of the world trade web. Scientific Reports **5**, 10595 (2015) https://doi.org/10.1038/srep10595

[51] Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate: A practical and powerful approach to multiple testing. Journal of the Royal Statistical Society. Series B (Methodological) **57**(1), 289–300 (1995). Full publication date: 1995

[52] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)

[53] Vallarano, N., Bruno, M., Marchese, E., Trapani, G., Saracco, F., Cimini, G., Zanon, M., Squartini, T.: Fast and scalable likelihood maximization for exponential random graph models with local constraints. Scientific Reports **11**(1), 15227 (2021) https://doi.org/10.1038/s41598-021-93830-4

[54] team, T.: Pandas-dev/pandas: Pandas. https://doi.org/10.5281/zenodo.3509134 . https://doi.org/10.5281/zenodo.3509134

[55] Harris, C.R., Millman, K.J., Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., Kern, R., Picus, M., Hoyer, S., Kerkwijk, M.H., Brett, M., Haldane, A., Río, J.F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., Oliphant, T.E.: Array programming with NumPy. Nature **585**(7825), 357–362 (2020) https://doi.org/10.1038/s41586-020-2649-2

[56] Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., Mulbregt, P., Contributors, S...: Scipy 1.0: fundamental algorithms for scientific computing in python. Nature Methods **17**, 261–272 (2020) https://doi.org/10.1038/s41592-019-0686-2

[57] Loper, E., Bird, S.: NLTK: The Natural Language Toolkit (2002). https://arxiv.org/abs/cs/0205028

[58] Hagberg, A.A., Schult, D.A., Swart, P.J.: Exploring network structure, dynamics, and function using networkx. In: Varoquaux, G., Vaught, T., Millman, J. (eds.) Proceedings of the 7th Python in Science Conference, Pasadena, CA USA, pp. 11–15 (2008)

[59] Csardi, G., Nepusz, T.: The igraph software package for complex network research. InterJournal **Complex Systems**, 1695 (2006)

[60] Peixoto, T.P.: The graph-tool python library. https://graph-tool.skewed.de (2014)

[61] Bastian, M., Heymann, S., Jacomy, M.: Gephi: An Open Source Software for Exploring and Manipulating Networks. http://www.aaai.org/ocs/index.php/ICWSM/09/paper/view/154

[62] Hunter, J.D.: Matplotlib: A 2d graphics environment. Computing in Science & Engineering **9**(3), 90–95 (2007) https://doi.org/10.1109/MCSE.2007.55

[63] Laken, P.P.: d3blocks: Python library for creating interactive visualizations using D3.js. https://github.com/d3blocks/d3blocks. Accessed: 2024-05-04 (2022)

[64] Inc., P.T.: Plotly: Collaborative data science. https://plotly.com/python/. Accessed: 2024-05-04 (2015)

[65] Buffa, L., Mazzilli, D., Piombo, R., Saracco, F., Cimini, G., Patelli, A.: Maximum entropy modeling of Optimal Transport: the sub-optimality regime and the transition from dense to sparse networks (2025). https://arxiv.org/abs/2504.10444

[66] OpenAI: GPT-4 Technical Report. https://openai.com/research/gpt-4. Accessed via OpenAI API (https://platform.openai.com) (2023)

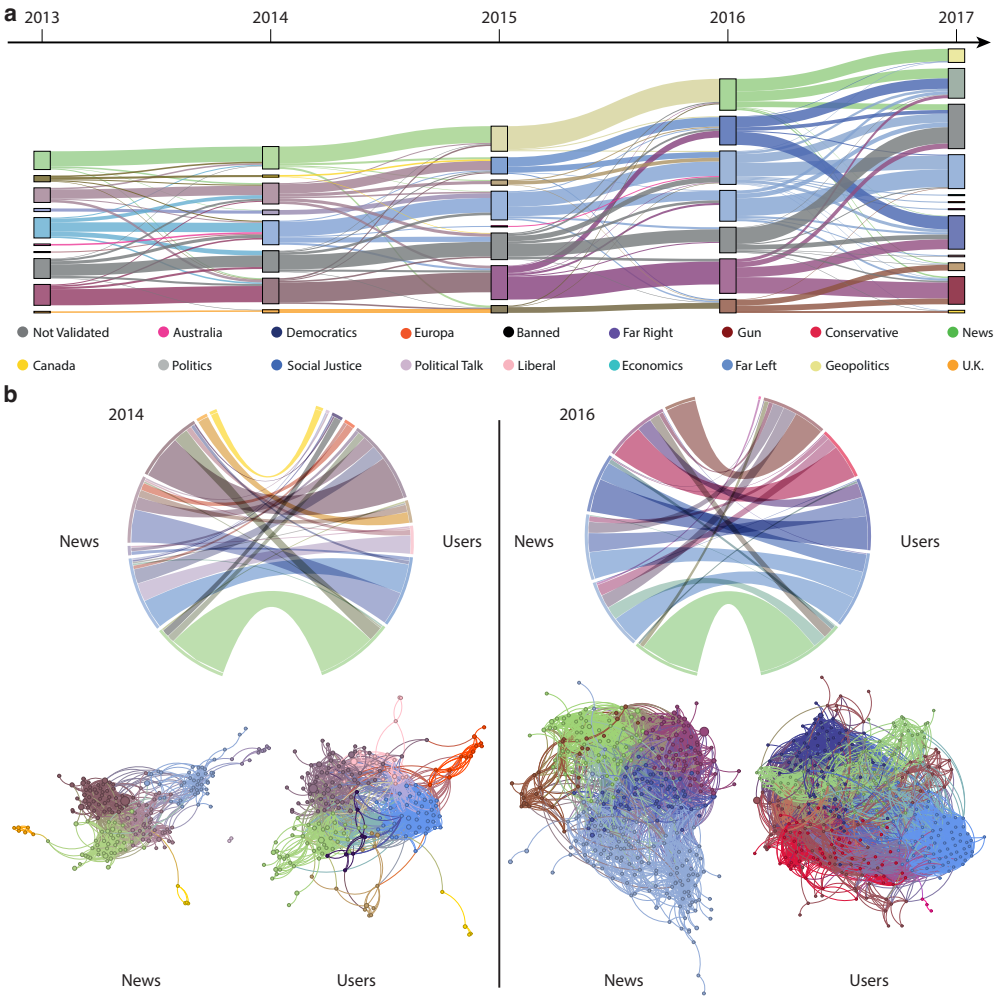

Fig. 4: **Patterns of News Consumption and Echo Chambers**. **a**: Sankey diagram where the width of the flows is proportional to the number of subreddits moving among communities, defined according to similarity of shared news sources among subreddits. Some communities remain highly homogeneous in terms of tags, while a merging of themes occurs in certain years. For example, Democrats and Conservatives initially coexist in the same community, but beginning with the election year, they gradually separate into distinct groups, eventually forming separate Democrats and Conservatives/Banned communities by 2017. **b**: Chord diagrams depicting the overlap (in terms of common subreddits) between communities identified through the "user-interaction" and "information-diet" analyses, for 2014 and 2016 respectively. The color of the flows represents the average of the colors of the source and destination communities. The coherent structure of the communities defined by the two approaches is indicative of the presence of echo chambers in political debate. Indeed the strongest overlaps tend to occur between communities with similar political or thematic orientations. Additionally, the figure shows the validated networks derived from the chord diagrams.

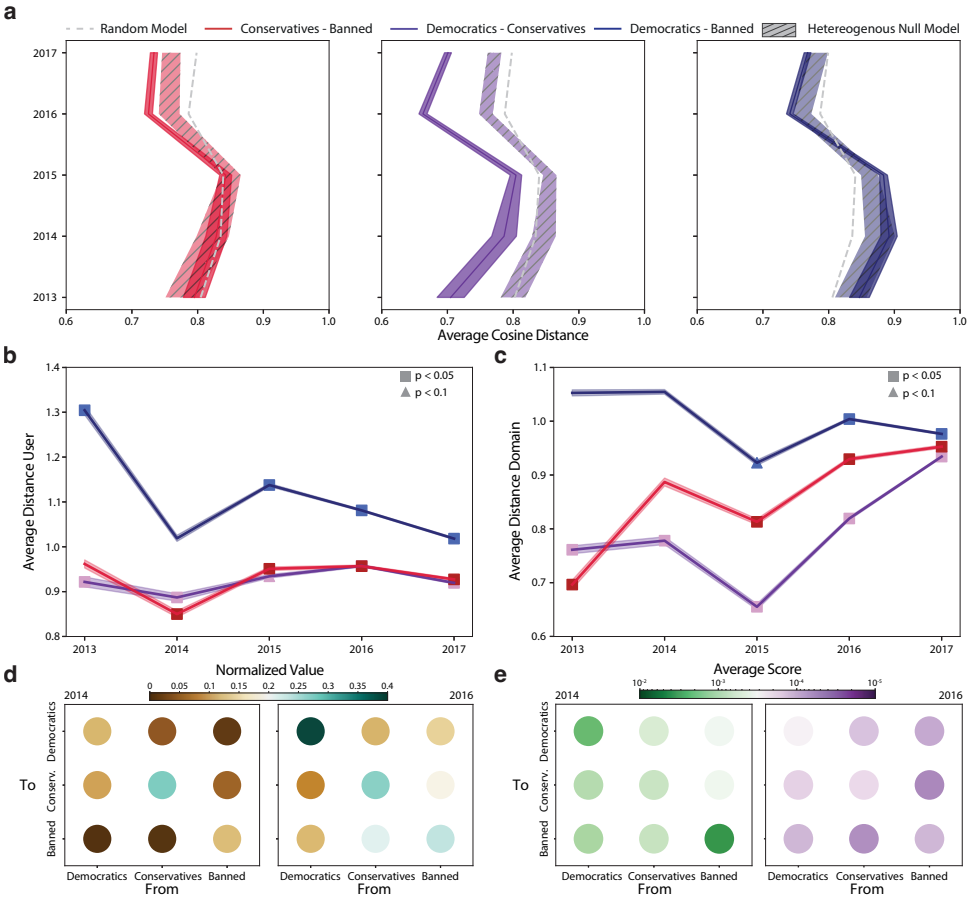

Fig. 5: **Engagement and Temporal Dynamics among Democratic, Conservative, and Banned Subreddits.**
a: Yearly textual distances between Democratic, Conservative, and Banned subreddits. We report yearly text-embedding cosine distances between Democrats and Conservatives, which decrease in 2016, particularly among politically active users. Both groups also converge linguistically toward Banned subreddits over time, although Democrats remain comparatively more distant. As benchmarks, we include a random model (the mean similarity of randomly assembled communities of the same size) and a heterogeneous model (the mean of all pairwise similarities among a community's constituent subreddits, thereby controlling for internal composition heterogeneity). **b**–**c**: Distances in user- and domain-based subreddit networks. We show average pairwise distances (normalized by yearly network averages) in statistically validated interaction networks, computed using the harmonic mean. A widening gap emerges between Democrats and Conservatives, with Banned subreddits consistently closer to Conservatives. Domain-sharing networks reveal a similar pattern: Democrats and Conservatives diverge sharply after 2016, while Banned subreddits remain aligned with Conservatives. **d**–**e**: Heatmaps of yearly proportions of comments and scores exchanged within and between user groups (shown for 2014 and 2016). They indicate increasing cross-faction interaction over time, but consistently lower comment scores for inter-group exchanges compared to intra-group ones. 26

Supplementary Files

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

- supplementary.pdf