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.

Anatomy of politically homogeneous spaces: How and why even they fragment

Ali Salloum

ali.salloum@aalto.fi

Department of Computer Science, Aalto University https://orcid.org/0000-0002-2381-6876

Dorian Quelle

Department of Mathematical Modeling and Machine Learning Digital Society Initiative, University of Zurich https://orcid.org/0000-0001-8257-7214

Letizia Iannucci

Department of Computer Science, Aalto University https://orcid.org/0000-0002-0635-9933

Titus Pünder

Department of Computer Science, Aalto University

Alexandre Bovet

Department of Mathematical Modeling and Machine Learning Digital Society Initiative, University of Zurich https://orcid.org/0000-0003-3937-3704

Mikko Kivelä

Department of Computer Science, Aalto University https://orcid.org/0000-0003-2049-1954

Article

Keywords: bluesky, political communication, polarization, social media

Posted Date: August 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10305465/v1

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

Additional Declarations: There is **NO** Competing Interest.

ANATOMY OF POLITICALLY HOMOGENEOUS SPACES: HOW AND WHY EVEN THEY FRAGMENT

Ali Salloum
Department of Computer Science
Aalto University
ali.salloum@aalto.fi

Dorian Quelle
Department of Mathematical Modeling and Machine Learning
Digital Society Initiative
University of Zurich
dorian.quelle@uzh.ch

Letizia Iannucci
Department of Computer Science
Aalto University
letizia.iannucci@aalto.fi

Titus Pünder
Department of Computer Science
Aalto University
titus.puender@aalto.fi

Alexandre Bovet
Department of Mathematical Modeling and Machine Learning
Digital Society Initiative
University of Zurich
alexandre.bovet@uzh.ch

Mikko Kivelä
Department of Computer Science
Aalto University
mikko.kivela@aalto.fi

July 9, 2026

ABSTRACT

We often frame political polarization as a binary struggle between two starkly opposed and homogeneous camps, yet our understanding of the internal pressures within these insulated in-groups remains surprisingly thin. Simplified perspectives often dismiss these in-groups as uniform, reciprocal 'echo chambers' or 'bubbles' where shared beliefs collapse into a singular, cohesive consensus. By analyzing a corpus of 100 million posts with their engagement patterns on Bluesky, we challenge the setting where all users within a system are assumed to share a completely unified perspective. We show that systems can reach a user base where over 97% of individuals share the same dominant political stance, and that this alignment does not automatically create an environment where ideas and opinions resonate in a uniform way. Instead, our results are consistent with the theory that less attitudinal tension allows the attention of the masses to splinter into specialized factions. Our findings show how political in-groups internally differentiate, suggesting that insulation relaxes the pressure for a united front and allows for the emergence of complex, multi-centered epistemic systems.

Keywords bluesky · political communication · polarization · social media

1 Introduction

It is widely thought that people have a tendency to form groups with other like-minded individuals [1] and are more likely to engage with beliefs, ideas, and opinions that resonate with their pre-existing views [2]. So when the internet started to gain more traction at the turn of the millennium, some scholars were concerned that it would increase *ideological polarization* and 'threaten democracy', because of people's new ability to 'filter out everything except what they wish to see, hear, and read' [3, 4, 5, 6]. However, years later, other researchers found evidence that such segregation on the internet was rather low due to the incidental exposure to cross-cutting views within diverse networks [7], bridging weak ties that reduce polarization [8], and heterogeneous online news consumption patterns [9].

As the internet evolved, attention turned toward emerging social media platforms. The digital transformation of the public sphere has established these platforms as important arenas for daily political discourse and debate, prompting researchers to use these environments as rich data sources for understanding contemporary political dynamics [10, 11, 12, 13, 14, 15].

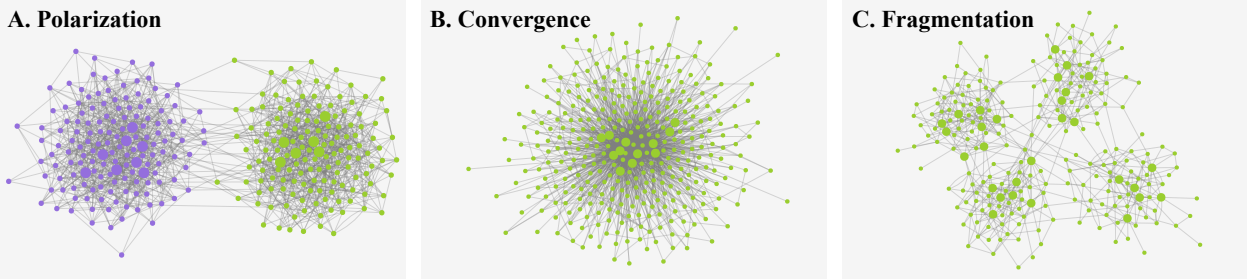


Figure 1: Beyond the standard polarization model. What happens when an in-group becomes ideologically insulated? Does it gravitate toward convergence, or develop internal fragmentation patterns? (A) *Polarization:* Represents the conventional view of 'echo chambers' or 'opinion groups' as two equal-sized, homogeneous, and assortative 'bubbles' characterized by strong in-group connectivity and out-group separation. (B) *Convergence:* Illustrates a maximally unified configuration where the insulated majority consolidates into a single, dense core-periphery structure, driven by pressures to maintain a united front and reinforce internal cohesion. (C) *Fragmentation:* Depicts the alternative, where the majority group internally fractionates into multiple smaller units. Rather than a uniform mass, this structure displays a division of intellectual labor, where subgroups can specialize in specific narratives, distinct aspects of an issue, or varying threat perceptions, creating a complex multi-core epistemic system.

For years, traditional research has focused on multi-partisan digital spaces (such as Facebook, Twitter/X) where opposing factions, leaders, and followers shared a singular environment, maintaining at least some degree of cross-group interaction and incidental exposure to out-group perspectives [13, 15, 16]. Recently, if anything, the digital ecosystem has seen a growing trend of users migrating into segregated, ideologically aligned spaces on both the political left and right [17, 18, 19, 20]. Some scholars characterize this transition as a shift from localized 'echo chambers' within shared environments to entirely insulated 'echo platforms' [17].

Many seminal works warned of 'disparate islands' and 'cyber ghettos' [12] of political communication more than two decades ago, viewing them as possible products of the extensive communicative heterogeneity enabled by the then-emerging internet [21, 12]. It is suggested that such systems may erode the integrative capacity of the public sphere and increase intolerance toward out-groups [12]. Yet, this same literature proposes reasons why these systems can coexist in a deliberative democracy. They can be viewed as a necessity for certain groups, acting as secure environments that allow communities to resolve internal conflicts and build a shared identity [12]. In some contexts, the prevailing political climate may prompt the emergence of a counter-public sphere [22], whereby groups can coordinate their efforts to facilitate social change or serve as a counter-force [22, 12].

In politically polarized settings, it is typical for both research and media commentary (e.g., [23]) to focus on the differences between oppositional camps, often deeming (either intentionally or inadvertently) the respective in-groups attitudinally and behaviorally uniform (see Fig. 1A) [24, 13, 25]. We also know that groups tend to develop highly reductive and stereotypical perceptions of the out-group as polarization intensifies [26]. But what actually occurs within these ideologically insulated enclaves? Does the absence of a prominent political out-group and/or a need for shared beliefs truly push the community toward a single, cohesive consensus, or does the environment allow for internal differentiation? On this question, competing theories and empirical findings point in opposite directions.

On one hand, insulation is often seen as a culmination of polarization. A wide body of work predicts that such environments should reinforce internal cohesion [17], intensify consensus [27, 4], and heighten collective opposition to external groups [26, 28]. These expectations draw on research in social identity and self-categorization, group polarization [27, 4], and selective exposure [29], which, in turn, describe distinct but complementary mechanisms underlying these effects. In fact, previous research shows that groups of like-minded people favor information that supports their views much more strongly than individuals do on their own [30].

On the other hand, a contrasting line of theory suggests that ideological insulation does not simply unify, but fragments the underlying social organization of the group. In the absence of direct engagement with an opposing out-group, the pressure to sustain a united front relaxes [31]. This, in turn, can lead actors to gravitate toward separate, sometimes competing, factions driven by either *emergent* or *intentional* causes. Emergent differentiation arises as a byproduct of *bounded rationality* [32, 33]. Because human attention and cognitive capacity are finite, individuals must selectively prioritize trusted sources, necessitating a natural *division of intellectual labor* to filter information [32]. This specialization is further reinforced by divergent sub-frames that elaborate on distinct aspects, problems and solutions within the same high-level cause [34]. Conversely, intentional differentiation develops from active social and political friction. Variations in threat perception create a spectrum of urgency [35], while conflicting attitudes toward authority can alienate establishment loyalists from internal critics [36]. Ultimately, while a shared baseline of values persists, these pressures can produce political schisms, in some cases even leading to the eventual rupture of the group [37, 38].

Given these two possible directions, the in-group can either show *convergence* or *fragmentation*. To investigate this, we conceptualize these digital environments as attention systems modeled as networks, where users represent nodes and their endorsements form the edges [39]. This approach allows us to operationalize a *maximally* centralized sphere as a state where every user's *endorsement footprint* converges onto the exact same set of central targets. In other words, if users across the system attend to a common set of leading voices, it should emerge as a structural configuration we term the *single gaze* (see Fig. 1B). This core-periphery structure matches prior findings on centralized online leader-follower hierarchies, in which a tight core of influential voices drives the discourse and the broader network functions as a uniform mass of amplifiers [40, 15].

Conversely, the contrasting perspective suggests that an insulated community is unlikely to show such structure. Instead of the system having a single gaze directed toward a centralized core, the group organizes into distinct, specialized factions with divergent focuses. In this scenario of a more fragmented sphere, the united front gives way to a *splintered gaze* (see Fig. 1C). Then the extent of this internal fragmentation can be measured by evaluating the divergence between individual endorsement footprints.

Conventional frameworks for analyzing polarization in social media cannot be directly applied to answer our question about ideologically insulated spaces. This is mostly because the majority of these methods are designed for social media discourse assumed to consist of exactly two completely opposing and homogeneous groups (see Fig. 1A) (e.g., [41, 13, 42, 43, 44, 39]). While some of these methods do recognize the internal leader-mass hierarchies within groups [15, 40, 13], they still assume that these hierarchies collapse into a uniform single gaze, omitting the possibility of more complex, multi-centered attention patterns. Consequently, this rigid two-camp paradigm breaks down in highly uniform spaces where a single political stance accounts for the vast, overwhelming majority of the user base, while completely opposing attitudes typically represent a tiny fraction.

We introduce a data-driven approach that, although applicable across diverse digital environments, is specifically designed to work for homogeneous spaces too. Our framework is built upon two fundamental requirements: the *organic* discovery of issue salience and *structure-agnostic* network inference (see Fig. 2). First, we adopt an unsupervised topic inference to identify issue salience, allowing us to find what issues a specific community actually prioritizes. This is necessary because researchers cannot know beforehand which themes matter to an enclosed population, making it risky to externally pre-define their discourse. Second, we apply network inference to uncover latent groupings of users. By letting the structural configurations emerge directly from the observed interaction data, we avoid forcing the network into artificial constraints, such as assuming that online groups must form assortative or binary structures.

In this paper, we turn to Bluesky, an emerging microblogging platform [45] whose rapid expansion was catalyzed by user migrations away from X following Elon Musk's acquisition [46, 19]. While past work has shown tentative evidence of a left-leaning political alignment on the platform (primarily through localized proxy metrics like domain sharing [19]), we substantiate this homogeneity in a more holistic and systematic manner. Starting from all platform content generated within our observation window (Jan–Dec 2025), we filtered the data to 99.97 million posts and classified them into twelve political issues. With this, we demonstrate that in ten of twelve issues, over 95% of active users share the dominant stance, with the underlying discourse mostly characterized by progressive and anti-conservative attitudes.

Despite this exceptionally high level of ideological uniformity, we show that these issue-specific systems consistently fragment. We find that issues rallying against an external enemy, however, lean toward structural convergence, whereas internal deliberation regarding core progressive priorities tends to fracture into decentralized factions.

2 Results

2.1 Bluesky hosts an active, diverse and critical public sphere

Bluesky has reached approximately 41 million registered users [47], with roughly 2 million individuals active on the platform daily. During our twelve-month observation window from January to December 2025, we recorded 1.6 billion posts and over 10 billion interactions. The platform exhibits substantial growth, averaging 46,000 new sign-ups per day. For this study, the final filtered dataset contains 99.97 million English-language posts that received at least one repost (see Methods 4.1).

Our unsupervised topic inference of this corpus (see Methods 4.2 and Fig. 2A) identifies 122 granular *microtopics*, revealing a discourse that is both intensely critical and vast in scope. Within the Trump-centric sphere, topics focusing on Project 2025, anti-MAGA sentiment, and 'anti-fascist' mobilization carry an overt oppositional tone. This alignment characterizes the broader dataset, as evidenced by topics dedicated to historical fascist parallels, reproductive rights advocacy, and LGBTQ+ and transgender rights. Institutional and economic critiques, such as anti-authoritarian protests, the Epstein document release, demands for wealth taxation, further signal the discourse's political intensity. The discourse can be characterized by a strong liberal-progressive and anti-conservative sentiment (see Appendix I).

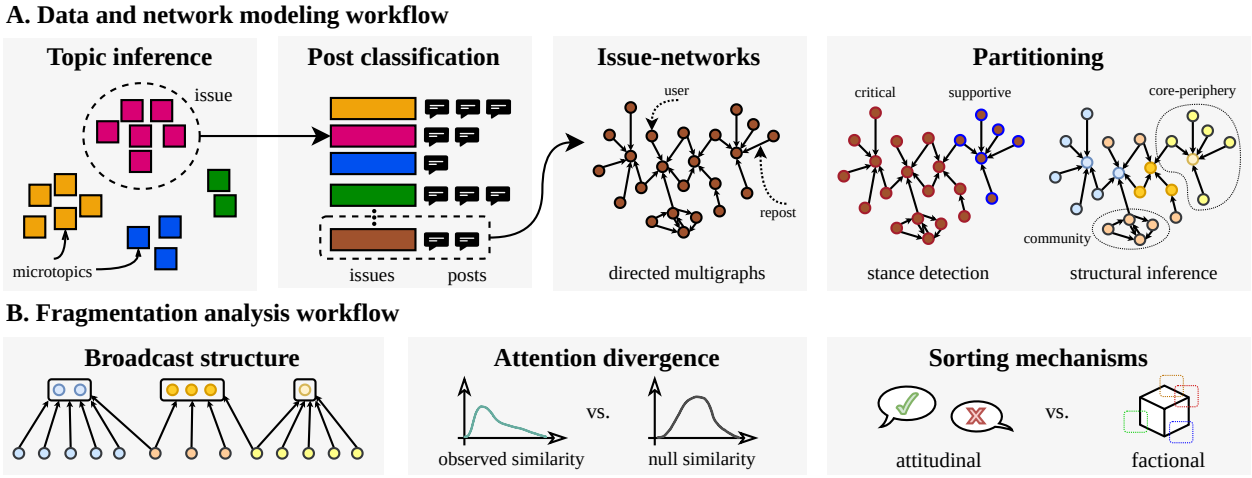


Figure 2: Data-driven pipeline for quantifying in-group fragmentation. The pipeline is divided into two main workflows: (A) Data and network modeling workflow and (B) Fragmentation analysis workflow. The process begins with *Topic inference* to identify granular microtopics, which are then used to construct the codebook for *Post classification*, where nearly 100 million posts are classified into twelve high-level political issues. These issues are then modeled as *Issue-networks* as directed multigraphs where nodes represent users and edges represent reposts. The workflow proceeds to *Partitioning*, which applies stance detection to classify users as critical or supportive and structural inference to identify structural-based grouping (e.g., core-peripheries, communities). In the second workflow, the *Broadcast structure* of these networks is analyzed to measure *Attention divergence*. This is done by comparing observed similarity in user endorsement footprints against a null similarity baseline that simulates a single gaze scenario. Finally, the pipeline assesses *Sorting mechanisms* to determine whether the observed fragmentation stems from attitudinal differences or factional sorting, wherein like-minded individuals divide their attention based on specialized interests and distinct topical niches.

The analysis also captures broader domestic and international conflicts. Within US politics, microtopics include opposition to Medicaid cuts, critiques of Supreme Court legitimacy, and the housing crisis. Critical discourse also emerges regarding internal party dynamics, specifically Democratic Party infighting and calls for Senate leadership changes. Global geopolitical discussions are represented by topics such as Ukrainian and Palestinian sovereignty, alongside shifts in the domestic politics of other nations, including UK political realignment and Mark Carney's rise.

Beyond political discourse, the data reveals a vast array of diverse subcultures and professional niches. These range from global entertainment trends like K-pop fandoms and indie game development to community welfare efforts such as animal rescue networks. Additionally, we find professional and lifestyle interests including AI research, wildlife photography, and explicit NSFW communities. For the full list of all 122 microtopics, see Appendix C.

2.2 About a fifth of the platform's content is political

These microtopics capture the organic variety of subjects discussed by users. We apply them as the empirical foundation for a taxonomy, which will be applied to map individual posts onto a structured discourse schema (see Methods 4.2). This bottom-up approach identifies twelve high-level issues, such as `Trump Administration`, `Russia-Ukraine Conflict`, and `Democratic Party (US)` (see Table 1).

To perform this classification across the entire corpus, we formalize these twelve issues into a comprehensive codebook. We write for each category a distinct definition and context (see Appendix D). This codebook serves as the system prompt for an ensemble of three independent, open-source LLMs used to classify the entire corpus of 99.97 million posts at the individual post level. To ensure the robustness of this classification, the codebook was rigorously iterated, and tested for sensitivity to guarantee that results remained stable and would not shift based on minor prompt variations. Posts inferred as political but not fitting these specific categories were labeled *Other Political*, while those outside the political sphere were labeled *Non-Political*. To aggregate model annotations while accounting for individual certainty, we developed a strategy called *Normalized Consensus Confidence*. This approach achieved over 90% accuracy, which was confirmed through a manual validation process conducted by human annotators (see Methods 4.3).

Nonpolitical interests account for the vast majority of activity on Bluesky, whereas political content represents at least 20% of the platform's reposted English-language posts. What stands out most in the data is the dominance of the `Trump Administration`, which alone accounts for 7.18% of these posts. To put its scale in perspective, users discuss it nearly five times as often as either of the global conflicts (1.27%–1.50%). We also find a substantial discourse around the US

Table 1: Categorization of the 99.97 million post dataset. These twelve issues, identified through the data-driven approach (see Methods 4.2), serve as the basis for classifying the dataset and constructing the political networks for all downstream analyses. We also report the prevalence of posts classified as other political, non-political, or ambiguous.

Category	High-Level Issues	Description	Post Count	Share (%)
US Politics	Trump Administration	Actions of US President Donald Trump and his admin.	7.18M	7.18
	Republican Party	Views on the US Republican Party.	1.42M	1.42
	Democratic Party	Views on the US Democratic Party.	670K	0.67
	Elon Musk	Political issues concerning Elon Musk.	754K	0.75
Intl. Politics	Russia–Ukraine	Ongoing conflict between Russia and Ukraine.	1.50M	1.50
	Israel–Palestine	Ongoing conflict between Israel and Palestine.	1.27M	1.27
Social Rights	Reproductive Rights	Legal and ethical status of abortion.	238K	0.24
	Gender & Sexuality	Rights and treatment of the LGBTQ+ community.	805K	0.81
	Equity & Rep.	Diversity, Equity, and Inclusion (DEI) programs.	196K	0.20
Systemic	Climate & Env.	Climate change and environmental policies.	745K	0.75
	AI Policy	AI regulation and societal impact.	480K	0.48
	Tech Platforms	Regulation of social media and tech platforms.	453K	0.45
Auxiliary	Other Political	Fallback class for other political or policy-related issues.	4.81M	4.81
	Non-Political	Examples include gaming, art, and explicit content.	72.61M	72.64
	Ambiguous	Instances with no majority consensus among classifiers.	6.83M	6.84
Total			**99.97M**	**100.00**

political parties, where nearly two times more posts were observed to be about the `Republican Party` at 1.42% than the `Democratic Party` at 0.67%.

We assess whether political discourse within these ideologically homogeneous environments converges into a single consensus by examining two key user characteristics. In Section 2.3, we infer each user's *high-level stance* (e.g., critical–supportive, for–against) from the content they engage with to capture their explicit positioning on an issue. Because public expressions can vary in scale or deviate from actual network behavior, in Section 2.4 we evaluate connectivity patterns to group users into *structural factions* based on their shared endorsement footprints. Finally, in Section 2.5, we cross-analyze these stances and structural factions to formally investigate the existence of a single gaze.

2.3 Bluesky is ideologically homogeneous

Platform hosts a liberal-progressive and anti-conservative sentiment

Prior work has found evidence of a large share of left-leaning link-sharing on Bluesky [19]. However, our approach shifts from proxy signals to direct classification of post content. By evaluating the actual text of all 99.97 million posts, we bypass the data sparsity and selection biases inherent to proxy-based methods, which are typically limited to indirect signals such as domain-level link analysis [48, 19], pre-selected keywords [11, 49], or hashtags [49] to infer user leanings. We find that Bluesky hosts a public sphere where political agendas and tone are closely aligned with a US-based liberal-progressive and anti-conservative consensus (see [50] and Appendix I), while showing less evidence for the traditional right-left or Republican-Democrat binary opposition.

To systematically quantify the scale of this homogeneity, we classify users across the twelve high-level political issues based on the content they authored or engaged with during our twelve-month observation window. Using situated labeling, a sampled collection of a user's issue-relevant content is categorized as 'Critical', 'Neutral', or 'Supportive' towards the issue in question (see Methods 4.4).

We observe strong ideological homogeneity across nearly all issues analyzed (see Table 2). In 10 of the 12 issues, the dominant stance accounts for more than 95% of the user base. This trend is consistent across categories, including political figures (`Trump Administration`, 99.50% critical; `Elon Musk`, 99.89% critical), international conflicts (`Russia-Ukraine Conflict`, 97.09% pro-Ukraine; `Israel-Palestine Conflict`, 97.55% pro-Palestine), and social policy (e.g., `Reproductive Rights`, 99.15% pro-choice). Consequently, minority opposition in these domains is negligible in size, typically falling below 2%. Our analysis also shows that, across all issues, the classification of users as 'Neutral' is rare, mostly ranging between 0.1%–2% of the population.

Table 2: Quantitative evidence of in-group homogeneity across issue networks. The table shows stance labels and user percentages per issue. Percentages are calculated only for users with at least five posts or reposts on the issue (see Appendix F for metrics using higher or lower thresholds). All AEI scores significantly exceed randomized null baselines ($p < 0.0001$), see Methods 4.5 for details. Note that AEI is not a percentage, but a score ranging from -1 to 1, where 0 indicates no polarization and 1 indicates maximum polarization between the dominant majority and minority stance groups. Negative values, which imply disassortative mixing, were not observed in our dataset.

	Majority Stance		**Minority Stance**		**Neutral / Other**	**Polarization**
Issue Network	*Label*	*%*	*Label*	*%*	*%*	*AEI*
`Trump Administration`	Critical	99.50	Supportive	0.21	0.29	0.99
`Elon Musk`	Critical	99.89	Supportive	0.04	0.06	0.99
`Republican Party`	Critical	99.74	Supportive	0.08	0.18	0.99
`Democratic Party`	Supportive	51.97	Critical	46.56	1.48	0.34
`Russia-Ukraine`	Pro-Ukraine	97.09	Pro-Russia	0.89	2.02	0.80
`Israel-Palestine`	Pro-Palestine	97.55	Pro-Israel	1.78	0.67	0.89
`Reproductive Rights`	Pro-Choice	99.15	Pro-Life	0.32	0.54	0.74
`Gender and Sexuality`	Supportive	98.05	Critical	1.79	0.15	0.45
`Equity and Representation`	Supportive	89.91	Critical	9.71	0.38	0.11
`Climate and Environmental`	Activist	98.25	Dismissive	0.37	1.39	0.40
`Artificial Intelligence`	Critical	95.90	Supportive	1.99	2.10	0.74
`Technology Platforms`	Critical	99.55	Supportive	0.30	0.15	0.59

`Democratic Party` stands out as the only exception to this pattern of attitudinal uniformity, representing the only truly divided domain *contentwise* in the dataset. We observe a substantial split between supporters (51.97%) and critics (46.56%). Both the emergent microtopics and manual inspection of the content confirm that this criticism does not imply these users are, for instance, Republican voters (see Appendix C). Instead, these critical users typically express 'disappointment' with current leadership or call for a 'significant shift' in the party's political approach (see Section 2.1).

Dissenting voices are rare and isolated

Despite the near-total attitudinal uniformity across most issues, we investigate whether the platform still exhibits the structural polarization patterns characteristic of more ideologically diverse social media environments. Specifically, we study if the interactional boundaries between the overwhelming majority and the marginal minority stances resemble the structural divides typically found in binary partisan systems (e.g., [13, 15, 39, 11]).

Modeling these specific discourses as *self-organizing social systems* enables us to analyze the interactional patterns that emerge among users, which has become a standard approach in studies of online polarization and echo chambers [39, 15, 41]. We follow this convention by transforming the classified posts into issue-specific networks through a projection onto the user dimension [39], i.e., in these networks, nodes represent individual users, and a directed edge exists from user i to user j if user i has reposted content authored by user j (see Methods 4.3).

The resulting issue-specific systems vary in size from approximately 367,000 to over 1.3 million nodes (see Table in Methods 4.3). Within this range, the `Trump Administration` network is the largest and most active, with an average of 62 reposts per user. For comparison, this system is approximately four times larger and eight times more active than the one representing `Reproductive Rights`, which represents the lower end of our observed political systems.

We compute the structural division between opposing camps using the *Adaptive External-Internal Index* (AEI) [41] (see Methods 4.5). This metric quantifies opinion polarization by measuring the extent of in-group favoritism and out-group avoidance, ranging from -1 to 1, where 1 indicates maximum separation. We observe that issues showing the highest ideological uniformity, specifically `Trump Admin.`, `Elon Musk` and the `Republican Party`, also display the strongest structural separation, with AEI scores reaching near-maximal levels (0.99). This indicates that the supportive minority, though microscopic in size (0.04%–0.08%), exists in a state of near-total isolation from the majority.

This structural behavior reveals a clear distinction between the impact on ideological adversaries and internal dissenters. While sharp interactional boundaries are drawn against out-group entities like Trump (and his administration), Elon Musk and Republican Party, the discourse of Democratic Party remains structurally cohesive. For it, we observe disagreement without fracturing. While the sentiment split is substantial (51.97% supportive vs. 46.56% critical), the `Democratic Party` exhibits a relatively low AEI score of 0.34. This suggests a softer separation with blurrier boundaries, where opposing factions are structurally entangled rather than isolated.

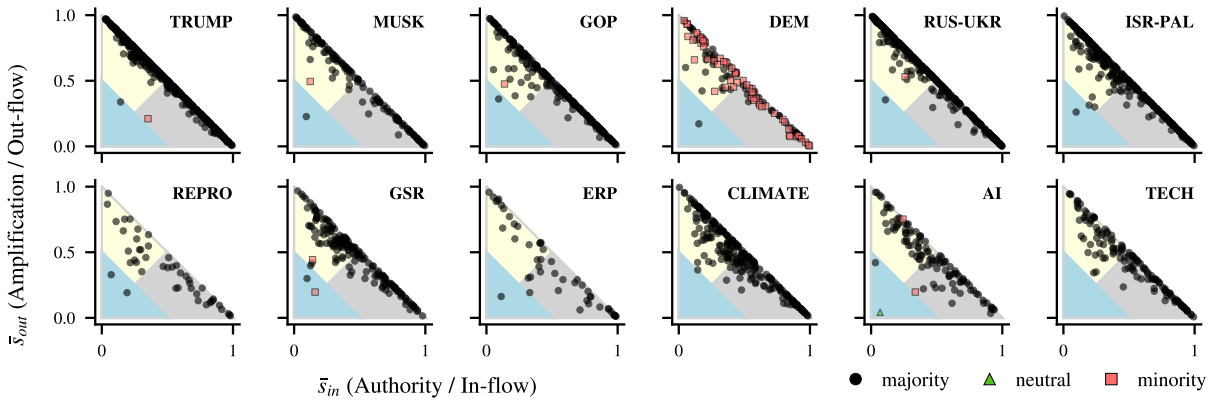


Figure 3: Multi-core-periphery topology dominates, while reciprocal communities remain rare. This figure displays the properties of blocks identified by the nested stochastic block model, mapped by their authority in-flow ($\bar{s}_{in}$) and amplification out-flow ($\bar{s}_{out}$). Because the sum of these flows is ≤ 1, all blocks are contained within the triangle bounded by the diagonal. It is evident that the majority of blocks have clear roles: they either primarily amplify other blocks (upper wing/yellow-shaded region) or receive attention as authorities (lower wing/gray-shaded region). There is little evidence of community-like blocks (inner triangle/blue-shaded area), where more than half of the reposts would be confined internally. Each block is categorized as majority, neutral, or minority based on the predominant stance of its members after the structural partition.

2.4 Users are organized into separate structural factions

To analyze the social organization of these networks, we set aside content-based labels and focus on the raw architecture of user interactions. We apply a nested degree-corrected stochastic block model (nested-SBM) [51] to identify the most probable partitions in the data based on statistical evidence of connectivity (see Methods 4.4). In doing so, we draw on the classical blockmodeling theory to infer social roles and macro structures purely from relational data [52].

After the structurally equivalent sets of users (the blocks) are uncovered, we assign them their roles within the system. We do this by characterizing every resulting block by two primary connectivity metrics [53]: the authority in-flow ($\bar{s}_{in}$), quantifying the volume of incoming attention, and the amplification out-flow ($\bar{s}_{out}$), quantifying the volume of attention projected outward (see Methods 4.6). By evaluating the balance between these two directional flows, we distinguish the authoritative core blocks (leaders) from the broad base of amplifying masses (see Methods 4.6). Furthermore, because this approach makes no prior structural assumptions, it also isolates assortative groups where empirically present.

Our analysis indicates that these issue networks are characterized by a *multi-core-periphery* structure (see Fig. 3). Instead of a singular core-periphery topology, the network features multiple independent core blocks acting as distinct targets of attention, while the remaining audience partitions into discrete peripheral blocks that function as localized amplifiers. Each peripheral block constitutes *a faction* characterized by a specialized endorsement footprint (see Methods 4.6), meaning that while these masses almost universally share an identical broad political stance, they are structurally differentiated by unique signatures of attention, i.e., their selective prioritization of different subsets of the system's leaders. This multi-centered division of attention contrasts with single-centered hierarchical models, which typically map the system onto a single elite and a uniform peripheral mass [40, 15].

While we observe that these networks are primarily self-organized for the amplification of preferred voices, some community-like groups emerge, albeit small and rare. For instance, in the `Trump Administration` and `Gender and Sexuality` systems, the only community-like structures to emerge are those primarily composed of minority supportive groups, which suggests that these ideological outliers rely on mutual reinforcement to sustain their presence within the network. Similarly, for `Artificial Intelligence` network, the only block exhibiting such internal circulation is a neutral community that, on manual inspection, shares AI research and publications among its members.

At this stage, each user is defined by their stance, structural faction, and role as either leader or mass. For the masses, these factors impact their endorsement footprints which give us the empirical basis to formally quantify the total magnitude and the quality of attentional convergence across our twelve systems (see Methods 4.6).

2.5 Homogeneous systems fragment differently for progressive policies and common adversaries

We begin by confirming that our issue-networks do fragment in the first place by measuring the distance (Φ) of the observed network from a maximally converged system (see Fig. 1B). For each system, we model the distribution of

Table 3: The fragmentation patterns of politically homogeneous systems. This matrix operationalizes the core theoretical mechanisms of in-group differentiation by mapping them directly onto our two empirical network metrics: Attitudinal Cleavage (ϕ_A) and Factional Divergence (ϕ_F).

	Low Factional Divergence (ϕ_F)	**High Factional Divergence (ϕ_F)**
Low Attitudinal Cleavage (ϕ_A)	**Convergence** The baseline scenario where the entire system converges on a singular, centralized set of core voices, leaving little room for mass factionalization, specialization, or a division of labor: 'echo platform'.	**Fragmentation** The shared attitudinal baseline relaxes the pressure for a unified front, allowing users to organically divide attention and specialize into distinct sub-issue tracks. Local opposing views are minimal or still share substantial overlap with the mainstream in attention patterns, which means there are still unifying voices in the system.
High Attitudinal Cleavage (ϕ_A)	**Polarized Convergence** The baseline convergence scenario but with a strongly isolated minority, where a dissenting group occupies a separate, detached block while the mainstream center remains highly unified around a central core: 'echo chamber in an echo platform'.	**Polarized Fragmentation** The system carries a strict boundary that excludes opposing perspectives, yet the insulated in-group still undergoes substantial internal splitting into specialized attention tracks.

pairwise similarity between users' endorsement footprints and compare this observed similarity against a null model (see Methods 4.7). Across all twelve issues, we find that Φ is consistently positive (see Appendix Table 11) and statistically significant (Mann-Whitney U, $p < 0.001$), rejecting the single gaze hypothesis in every domain.

To understand how this fragmentation is structurally organized, we break it down into two distinct sorting principles. The first is *Attitudinal Cleavage* ϕ_A, which tracks how much user attention splits between the stance majority and minorities (see Methods 4.7). The second is *Factional Divergence* ϕ_F, which measures the attentional differences within and between the distinct mass factions found in Section 2.4. Both metrics are calculated directly from the distribution shifts shown in our empirical cumulative distribution functions (see Fig. 4A), which capture the underlying pairwise similarity variations.

By plotting these two metrics against each other (see Fig. 4B), we can map our empirical findings directly onto the theoretical configurations outlined in Table 3. A baseline scenario with near-zero values across both dimensions would indicate a perfect single gaze. While this absolute *Convergence* regime remains completely empty in our data, the network for `Equity and Representation` serves as its closest empirical approximation ($\phi_A = 0.10$, $\phi_F = 0.22$).

We find that Attitudinal Cleavage is highest in systems centered on external antagonists, as seen in the `Republican Party` ($\phi_A = 0.55$) and the `Trump Administration` ($\phi_A = 0.52$). Yet, while these adversarial issues exhibit clear attitudinal divisions, their internal Factional Divergence remains low (`Trump Administration`: $\phi_F = 0.19$; `Elon Musk`: $\phi_F = 0.23$; `Republican Party`: $\phi_F = 0.25$). This combination indicates that for communities organizing around a common enemy, the differentiation into factions is much less observed compared to the remaining issues. Instead, they hold overlapping attention patterns. Empirically, these three issues occupy the *Polarized Convergence* quadrant of our matrix, characterizing a regime of a unified mainstream center paired with a strongly isolated minority.

In contrast, an inverse pattern emerges within core policy and ideological domains. The largest values for Factional Divergence are observed for `Artificial Intelligence` ($\phi_F = 0.43$), `Gender and Sexuality Rights` ($\phi_F = 0.40$), and the `Israel-Palestine Conflict` ($\phi_F = 0.37$). However, because these dimensions represent a continuous spectrum, these systems shift across different regimes based on their degree of attitudinal tension. For issues like `Gender and Sexuality Rights` ($\phi_A = 0.07$) and `Climate and Environment` ($\phi_A = 0.04$), high internal splitting occurs alongside very low attitudinal cleavages. This configuration maps onto the non-polarized *Fragmentation* state, where the in-group becomes divided into specialized or competing factions, and those with opposite stances still find something in common with the dominant group.

Conversely, `Artificial Intelligence` stands out as the closest approximation of *Polarized Fragmentation* in our dataset. The domain exhibits a pronounced Attitudinal Cleavage ($\phi_A = 0.34$), reflecting a far sharper attentional boundary between stance groups than in the core progressive issues, alongside the dataset's highest Factional Divergence ($\phi_F = 0.43$). Because the stance minority comprises under 2% of users, this divergence is overwhelmingly driven by pairs of like-minded users: even the critical majority itself splinters into specialized, independent attention tracks.

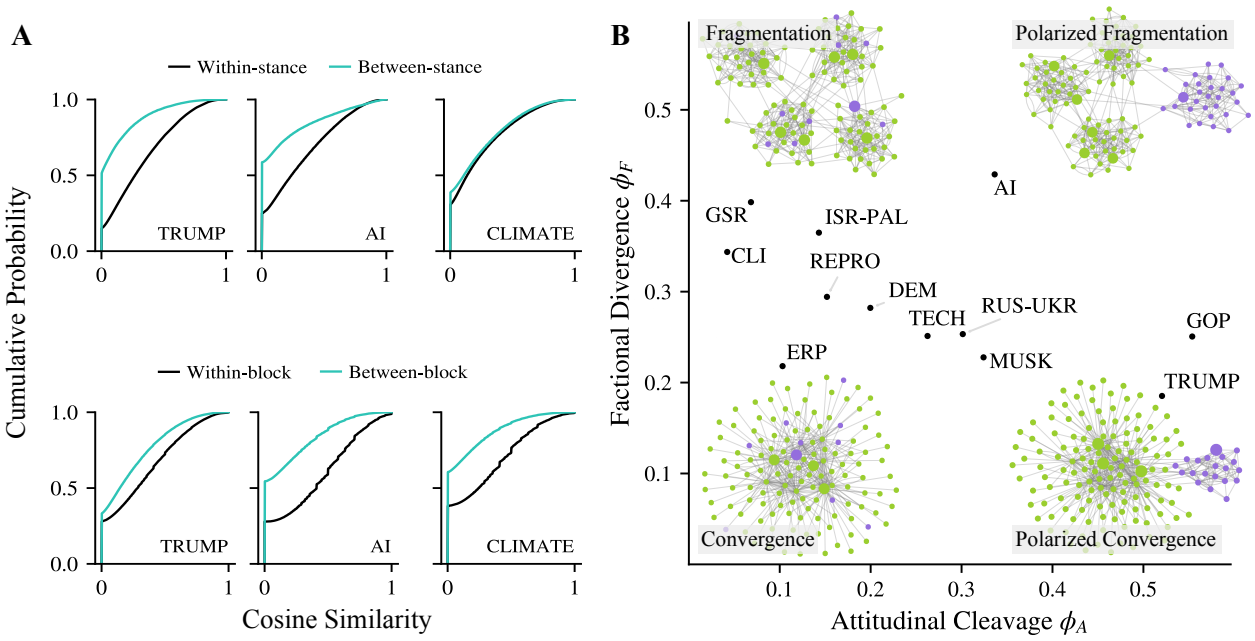


Figure 4: Quantifying mechanisms of in-group fragmentation and attentional divergence. (A) Empirical cumulative distribution functions of pairwise attention similarity for three selected issue networks (`TRUMP`, `AI`, and `CLIMATE`). The top row evaluates sorting by ideological stance (within-stance vs. between-stance), while the bottom row evaluates sorting by structural faction (within-block vs. between-block). A leftward shift of the between-group distributions (teal) relative to the within-group distributions (black) indicates lower attentional similarity across boundaries, reflecting the respective strength of each sorting mechanism. (B) Two-mechanism decomposition projecting the issue networks into a shared space of sorting principles: Attitudinal Cleavage (ϕ_A, x-axis) and Factional Divergence (ϕ_F, y-axis). The spatial layout maps out distinct fragmentation regimes: (1) core progressive priorities (e.g., `GSR`, `CLI`, `REPRO`) display non-polarized *Fragmentation*, in which attention splits across factions despite a shared value baseline; (2) external adversarial targets (e.g., `TRUMP`, `GOP`) show *Polarized Convergence*, characterized by a tightly consolidated mainstream majority paired with a deeply isolated minority fringe; and (3) emerging sociotechnical debate centered on `AI` is the closest approximation of *Polarized Fragmentation*, where the structural isolation of users with minority attitudes becomes much more pronounced than the fragmentation found among like-minded peers in the majority. The regime corresponding to the absolute *Convergence* remains empty. Insets visualize the characteristic network topologies that match with each quadrant regime.

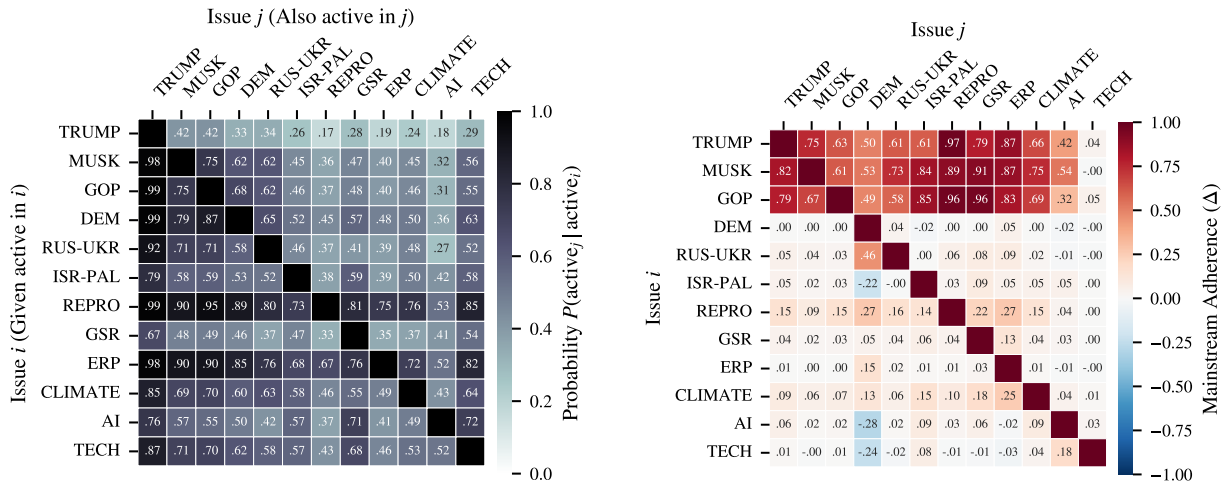
Our qualitative extraction (see Supplementary Information) from the factions of each issue provides illustrative contextualization for these distinct fragmentation regimes. In the `Climate and Environment` network, attention splinters across a wide array of specialized tracks, including clean energy transitions, localized conservation, plant-based dietary imperatives, and indoor air quality, all while maintaining a baseline attitudinal consensus. Slightly more contentious subniches, such as pro-nuclear advocacy or wildlife protection debates, coexist within the ecosystem without increasing the attitudinal tension too much.

The internal fracturing within `Artificial Intelligence`, by contrast, follows starker lines. Several critical factions evaluate the technology in strongly normative terms, focusing variously on pervasive surveillance, the spread of disinformation, copyright infringement, and the systemic theft of creative labor. These factions remain sharply separated from the more optimistic or neutral tracks, which center on the economic democratization of technology, geopolitical tech rivalries, and the falling cost of building top-tier systems such as DeepSeek.

Points of tension crack the united front

Until now, our analysis has treated each issue system separately. We now examine them jointly to assess how far the united front extends across issues. Two well-established accounts predict strong cross-issue alignment in such a setting: political positions tend to align across issues for the sake of *ideological coherence* [41], and where a majority overwhelmingly dominates, dissent may additionally be suppressed in line with the *spiral of silence* [54]. We test how far this alignment actually holds.

We first quantify audience overlap as the conditional probability $P(\text{active}_j \mid \text{active}_i) = |U_j \cap U_i|/|U_i|$ that a user active in issue i is also active in issue j (see Fig. 5A). Overlap is moderate to substantial but markedly asymmetric: users active in almost any issue are highly likely to also discuss the `Trump Administration`, while activity in the



(A) Activity Overlap

(B) Mainstream Adherence

Figure 5: Quantifying inter-issue alignment and tensions. (A) Activity Overlap: the conditional probability $P(\text{active}_j \mid \text{active}_i)$ that a user active in issue i (rows) is also active in issue j (columns). The matrix is asymmetric when issue sizes differ: a small issue's users may largely overlap with a large issue's audience while the reverse need not hold. (B) Mainstream Adherence, $\Delta[i,j] = P(\text{majority}_j \mid \text{majority}_i) - P(\text{majority}_j \mid \text{minority}_i)$. Positive values (red) indicate that issue i's majority adheres to issue j's mainstream more than issue i's minority does; negative values (blue) mark points of tension, where the majority is *less* likely to hold the mainstream stance on issue j than the minority; values near zero (white) indicate independence.

Trump discourse rarely predicts activity elsewhere. A block of issues around the `Trump Administration`, `Elon Musk`, the two US parties, and the `Russia-Ukraine Conflict` shares much of its audience (overlaps mostly above 0.60), whereas `Artificial Intelligence` and `Reproductive Rights` draw narrower segments of the platform. Some specific pairs, such as `Russia-Ukraine` and `Elon Musk`, or `Israel-Palestine` and `Gender & Sexuality Rights`, are more strongly connected than the two geopolitical conflicts are to each other.

Stances, too, align across issues, but only partially. Minority membership in one issue predicts minority membership in some others, measured by $P(\text{minority}_j \mid \text{minority}_i)$ (see Appendix L): supporters of the `Trump Administration` are disproportionately critical of the `Democratic Party` and pro-life on `Reproductive Rights`, yet remain predominantly pro-Ukraine on `Russia-Ukraine`. We therefore find no single, persistent dissident bloc that opposes the majority on every front.

For the majorities, we define *Mainstream Adherence*, $\Delta[i,j] = P(\text{majority}_j \mid \text{majority}_i) - P(\text{majority}_j \mid \text{minority}_i)$ (see Fig. 5B). In a *fully aligned* environment, Δ should not be be negative: holding the mainstream position on one issue should predict following the mainstream on any other at least as strongly as dissenting does. Negative values therefore flag *points of tension*, where a mainstream position on issue i is associated with a shift away from the consensus on issue j.

Such tensions appear most clearly relative to the `Democratic Party`: users holding the majority positions on `Israel-Palestine` (Pro-Palestine), `Artificial Intelligence` (Critical), or `Technology Platforms` (Critical) are less likely to support the party than the respective minorities are (permutation test, $p < 0.001$; see Methods 4.8). These instances hint at subtle friction in the platform's broad political alignment, as holding the mainstream view on one issue can weaken, rather than reinforce, adherence to the consensus on another, cracking what would otherwise be a united front *across* issues.

3 Concluding discussion

When polarization divides people into separate camps, these insulated groups do not necessarily collapse into a state where everyone agrees or attends to the same aspects of an issue. In this paper, we provide evidence for this by first demonstrating that Bluesky is politically homogeneous at the level of broad stances, and then analyzing the divergent attention patterns found among these like-minded users. We argue that these in-groups should be viewed as systems where users share core beliefs, while the distribution of attention across specific aspects of an issue remains heterogeneous.

The clearest evidence for this view is an absence. The configuration that popular intuition would predict for an 'echo platform', a maximally converged system in which a homogeneous mass attends to a single set of voices, never appears in our data: none of the twelve issue systems occupies the *Convergence* regime. Attitudinal consensus and attentional convergence are, in other words, distinct quantities, and they can move in opposite directions: the issues with the broadest agreement often exhibit the most divided attention.

Beyond the empirical case, this work contributes a general framework for a setting that existing methods struggle to handle. Conventional polarization measures presuppose two comparably sized, opposing camps and become degenerate when a single stance covers over 97% of users. Our pipeline — organic issue discovery, structure-agnostic network inference, and a null model that operationalizes maximal convergence under finite user activity — requires no such assumption, and the decomposition into attitudinal cleavage and factional divergence applies to any attention system, one-sided or not. Because it operates on a near-complete record of the platform rather than a sample, the framework also demonstrates that post-level ideological mapping at the scale of an entire social network is now feasible.

Simply dismissing these emerging homogeneous spaces as detrimental to discursive democracy, or writing them off as predictable echo chambers full of outraged partisans, impoverishes our understanding of contemporary political communication ecosystems. As noted earlier, much seminal work in political science and epistemology predicts the formation of such spaces for certain functions, such as identity building and division of labor. Their evolution is not inherently an alarming phenomenon.

What we find instead is that these homogeneous systems absorb users' divergent interests, disagreements, and less popular stances in ways that leave clear signs in the underlying structures. These signatures track what an issue is fundamentally about: whether the discourse concerns an internal priority, an external adversary, or an emerging sociotechnical debate.

In discussions surrounding core *in-group* priorities, such as climate politics, gender and sexuality rights, or internal debates about the Democratic Party, a divergence in focus does not result in hostile fracturing. This is consistent with theoretical accounts in which shared ideological alignment relaxes the pressure to sustain a united front [31]: user engagement flows across a continuous spectrum of attentional patterns, with even marginal or critical voices remaining structurally embedded within the broader network rather than isolated. The `Democratic Party` offers the sharpest illustration. It hosts the only genuinely contested stance distribution in our data, with supporters and critics near parity, yet the two sides remain structurally entangled rather than sorted into opposing blocks — disagreement without fracture. Inside this group of issues, many distinct factions coexist without producing complete epistemic silos.

For issues centered on adversarial figures and *out-groups*, interaction boundaries are more absolute. Between critics and supporters of Trump or the GOP, there is almost no shared ground or overlapping content. Within these domains, critics comprise over 99% of the user base, and the structural factions they form stay close to one another: they share a common focus on defensive and oppositional voices, and remain far more strongly tied to each other than the factions in any other political topic.

These absolute boundaries between stance groups begin to blur when the conversation moves to sociotechnical issues, such as AI policy and tech platforms — and, in part, Elon Musk, whose discourse is structurally adversarial but whose attitudinal boundary is markedly softer than Trump's or the GOP's. These systems show a higher attitudinal cleavage than core in-group priorities, but their attention networks are clearly less separated than those surrounding strictly political actors. One plausible explanation for this softening is the relative novelty and cross-cutting nature of these issues, which have not *yet* hardened into ideological identity markers. An individual can be excited by AI developments or space exploration, but also concerned that legislation is failing to keep pace or that these systems are heading in a troubling direction. Since these issues do not map onto traditional left–right divides, users can hold such seemingly conflicting views at once. Moreover, the level of urgency varies across communities depending on their specific professional priorities. How these sociotechnical issues will evolve, and whether they will ultimately drift toward traditional political lines, remains to be seen.

We also show that the presence of dominant majorities does not automatically guarantee cross-issue unity: specific mainstream positions coincide with measurable drops in adherence to the platform consensus elsewhere. We see this notably when contrasting the majorities in AI, tech regulation, and Israel–Palestine with the mainstream supportive of the Democratic Party. Even though attitudes critical of big tech and supportive of Palestine represent the dominant consensus on those issues, these users show a pronounced drop in adherence to the Democratic establishment. While we cannot say whether these tensions are persistent, they currently mark the clearest fault lines for potential political schisms within that party.

Our observations of how attitudinally homogeneous communities organize themselves in different ways, and of how local minorities either blend into the majority or become isolated, challenge the popular caricatures of 'echo platforms' that dominate public commentary. Instead of falling into flat uniformity, individuals within this sphere remain attracted

to specific ideas, voices, and figures with a fluidity that shifts depending on the nature of the domain. Acknowledging this internal complexity allows a more accurate view of how politically homogeneous environments organize, coalesce, and fragment.

Our study is, of course, confined to a single platform, a single year, and repost-based attention. Whether the tensions we detect will fade or deepen, and whether insulated platforms on the political right fragment along similar lines, are questions this work cannot answer, but ones the framework is built to ask. As users continue to sort into ideologically aligned spaces, these questions will only grow in relevance: internally differentiated homogeneity may prove to be the default condition of online political life.

4 Methods

4.1 Data

Collection For this study, we collected the complete stream of public network data from the Bluesky social network spanning the period from 1 January 2025 to 31 December 2025 using the AT Protocol Firehose API [55]. Unlike restricted sampling methods often required for established platforms, this approach allowed us to ingest the full firehose of user activity in real-time. Our data collection captured the activity of every user who performed at least one visible interaction during the observation window, covering all supported languages. We maintained a collection uptime of over 99.2% throughout the twelve-month period. Details on the ingestion architecture are provided in Appendix A.

Filtering We filtered the raw dataset to include only posts that had received at least one repost and contained a minimum of fifteen characters. Furthermore, because our preliminary analysis indicated that over 80% of the platform's content was in English, we restricted this study to English-language posts using Bluesky's native language metadata. After filtering, the dataset contained 99.97 million posts.

4.2 Topic inference and issue taxonomy

Microtopic identification To provide a data-driven and unbiased overview of the discourse taking place on the platform, we developed a fully computational pipeline that transforms raw posts into semantically meaningful clusters, which constitute our *microtopics*. Using a 10% sample of the 99.97 million posts, each post was embedded into a dense 248-dimensional vector with `Qwen3-Embedding-8B` [56], projected into a 12-dimensional space with UMAP [57], and grouped into discrete clusters with HDBSCAN [58], which identifies dense regions of similar posts. Pipeline parameters were calibrated through a grid search, with the final configuration selected on the *Silhouette* score [59]. To interpret the resulting clusters, we report for every microtopic the top ten TF-IDF keywords together with an LLM-generated title and thematic description, based on a representative sample of 100 posts per cluster. Full methodological details, including the hyperparameter grid and the selected parameters, are provided in Appendix B.

Aggregation into high-level issues Guided by these keywords and summaries, we manually aggregated the 122 granular microtopics into twelve broader political issues, such as `Trump Administration`, `Israel-Palestine Conflict`, and `Democratic Party (US)` (see Table 1). This bottom-up process ensures that the issues are empirically grounded in the platform's actual discourse rather than imposed a priori, and they serve as the target taxonomy for classifying the full dataset in all subsequent analyses. The full categorization of both political and non-political microtopics, along with their assignment to specific high-level issues, is provided in Appendix C.

4.3 Building the political systems

Issue classification We formalized the twelve issues into a codebook (see Appendix D), which served as the system prompt for an ensemble of three open-source LLMs from independent developers, chosen to reduce correlated model biases: `Mistral-Small-4-119B-2603-NVFP4`, `Gemma-4-26B-A4B-it`, and `Qwen3-30B-A3B-Instruct-2507`. The ensemble classified the entire corpus of 99.97 million posts into one of the twelve issues or three auxiliary labels: *Other Political*, *Non-Political*, and *Ambiguous* (a full disagreement among the three models). The codebook was iteratively refined and tested for sensitivity to ensure that classifications remained stable under minor prompt variations. To aggregate the models' outputs, we evaluated multiple voting strategies and found a consensus model based on *Normalized Consensus Confidence* (NCC) to be Pareto-optimal, achieving over 90% accuracy at 90% coverage when validated against a random sample of 1,000 posts manually labeled by experienced annotators (see Appendix E). Detailed benchmarks of the tested strategies, including precision–recall trade-offs and error analyses, are provided in the Supplementary Information.

Network construction Let X denote the set of twelve high-level issues and, for each $x \in X$, let P_x be the set of posts assigned to x. We define V_x as the set of users who authored or reposted at least one post in P_x during the observation window (January–December 2025). For each issue, we construct a directed multigraph $G_x = (V_x, E_x)$: every repost by user i of a post $p \in P_x$ authored by user $j \neq i$ contributes one directed edge (i, j) to E_x. Parallel edges are preserved, so the multiplicity m_{ij} counts how many of j's issue-x posts i endorsed. Self-reposts are excluded, as are quote posts, since the latter are frequently used to contextualize or criticize content rather than endorse it. All analyses in this paper operate on these twelve aggregated repost multigraphs. The construction extends directly to other interaction types (likes, follows, blocks) and to time-resolved snapshots $G_{x,t}$. The constructed networks vary in size and density across the twelve issues, with node counts of at least 360,000 users and the vast majority of users belonging to the largest weakly connected component. Full topological statistics for all twelve issue-specific networks are detailed in Table 4.

Table 4: Descriptive statistics of the issue-specific political networks. The table reports the number of unique users ($|V|$), reposts ($|E|$), average degree ($\langle k \rangle$), and the percentage of nodes belonging to the largest weakly connected component (% LWCC) for each of the twelve constructed graphs. See the Supplementary Material for further network properties and degree distributions.

Issue Network	Nodes ($\|V\|$)	Edges ($\|E\|$)	Avg. Deg. ($\langle k \rangle$)	% in LWCC
`Trump Administration`	1,337,239	83,005,561	62.1	99.4%
`Elon Musk`	711,898	10,406,742	14.6	99.3%
`Republican Party`	660,781	13,404,344	20.3	99.4%
`Democratic Party`	549,769	7,965,948	14.5	99.4%
`Russia-Ukraine Conflict`	605,423	14,585,247	24.1	99.1%
`Israel-Palestine Conflict`	566,748	8,477,056	15.0	99.1%
`Reproductive Rights`	366,266	2,574,102	7.0	98.6%
`Gender and Sexuality Rights`	770,675	8,871,379	11.5	98.7%
`Equity and Representation`	442,726	2,666,026	6.0	98.4%
`Climate and Environment`	548,278	5,515,629	10.1	98.6%
`Artificial Intelligence`	519,468	3,716,558	7.2	97.9%
`Technology Platforms`	686,161	5,595,502	8.2	98.6%

4.4 Stance detection and structural partitioning

Content-based stance detection To determine the ideological positions of users within each issue-specific network, we sampled for every user up to $n = 20$ of their posts and reposts relevant to the issue (or their entire relevant history, if shorter). These samples were analyzed by the ensemble classifier described above, adapted for stance detection through issue-specific system prompts that map each user's content profile to a discrete stance category tailored to that domain.

Labels were selected to capture the polarized binaries through which an in-group itself perceives these issues. While we use evaluative labels (e.g., 'Critical' vs. 'Supportive') for political entities and wars, for emerging sociotechnical issues we employ context-aware, or *situated*, labeling [60, 61]: labels reflect the normative framing dominant within this specific partisan environment, predominantly a US-based liberal-progressive consensus. We define stances relative to this in-group consensus, acknowledging that such positions are often constructed in opposition to an ideological 'Other', that out-group members would characterize their own positions differently, and that reducing nuanced political discourse to discrete labels necessarily simplifies. Our validation nevertheless indicates that these categories capture the general leaning of users well; full label definitions, prompt templates, and evaluation metrics are provided in the Supplementary Information.

Structure-based group detection We analyzed the structural organization of the networks using the *nested degree-corrected stochastic block model* [51]. Unlike modularity-based methods that optimize for internal density, the SBM is a probabilistic generative model that partitions the network into blocks of nodes that are *stochastically equivalent*. This distinction is essential for our analysis: the SBM assumes no specific mixing pattern such as assortativity, remaining structure-agnostic and capable of uncovering dense communities, disassortative hierarchies, or core-periphery structures, depending only on the statistical evidence in the data. Model inference follows the *minimum description length* principle, which balances the quality of the fit against the complexity of the model [62]. The nested variant additionally overcomes the resolution limit inherent in standard community detection, allowing us to resolve small structural units, such as minority sub-communities, that would otherwise remain undetectable in networks of this scale.

We select Level 2 of the inferred nested partition to balance structural granularity with statistical stability. To check that our topological findings are robust across scales, we repeated the analysis at Levels 1 and 3; these results, which confirm the consistency of the observed broadcast architecture, are provided in Appendix H.

4.5 Structural polarization index (AEI)

To quantify the structural division between the majority and minority stance groups within each issue-specific network, we utilize the *Adaptive External-Internal Index* [41, 39, 15]. This metric assesses the extent of in-group favoritism and out-group avoidance by comparing the density of interactions within stance groups to the density between them:

$$P_{\mathrm{AEI}} = \frac{i_\alpha + i_\beta - 2e_{\alpha\beta}}{i_\alpha + i_\beta + 2e_{\alpha\beta}}, \tag{1}$$

where i_α and i_β denote the internal link densities within groups α (majority) and β (minority), and $e_{\alpha\beta}$ the external link density between them. The index ranges from -1 to 1, with 1 indicating maximum structural separation. To confirm that the observed divisions are not artifacts of the network's basic features [39], we tested our results against a null model in which edges were randomly shuffled while preserving the degree sequence; observed polarization significantly exceeds this random baseline for every network ($p < 0.0001$).

4.6 Block roles and attention profiles

Since the nested stochastic block model infers the most probable partition directly from the data, the resulting structures cannot be known a priori. To map them, we adopt the diagrammatic approach of Hric et al. [53], which characterizes each block r by the relative proportions of its connection patterns: incoming reposts ($\bar{s}_{\mathrm{in}}^{(r)}$), outgoing reposts ($\bar{s}_{\mathrm{out}}^{(r)}$), and internal reposts ($\bar{s}_{\mathrm{int}}^{(r)}$), each normalized by the block's total repost activity. Because the three fractions sum to one, plotting $\bar{s}_{\mathrm{in}}^{(r)}$ against $\bar{s}_{\mathrm{out}}^{(r)}$ places every block inside a triangle whose regions correspond to distinct structural roles: blocks in the inner triangle are community-like in a weak sense (more than half of their repost activity is internal), blocks in the upper wing primarily amplify other blocks, and blocks in the lower wing primarily receive attention as authorities.

To separate the amplifying masses from the authoritative leaders, we collapse these flows into a single role index for each block r:

$$a_r = \frac{\bar{s}_{\mathrm{in}}^{(r)} - \bar{s}_{\mathrm{out}}^{(r)}}{\bar{s}_{\mathrm{in}}^{(r)} + \bar{s}_{\mathrm{out}}^{(r)}}, \tag{2}$$

which maps each block onto a continuous spectrum ranging from pure amplification ($a_r \approx -1$) to dominant authority ($a_r \approx 1$). Denoting the set of all inferred blocks by B, we partition it into the amplifying mass blocks $M = \{r \in B \mid a_r < -\frac{1}{3}\}$ and the authoritative core blocks $C = \{r \in B \mid a_r \geq -\frac{1}{3}\}$. The threshold $a_r < -\frac{1}{3}$ is equivalent to requiring $\bar{s}_{\mathrm{out}}^{(r)} > 2\,\bar{s}_{\mathrm{in}}^{(r)}$, ensuring that mass blocks act primarily as disseminators rather than originators of content.

Because the vast majority of users belong to the amplifying blocks, we aggregate them into a single pool, $U_{\mathrm{mass}} = \bigcup_{r \in M} V_r$, where V_r denotes the set of users in block r. The $k = |C|$ authoritative blocks are preserved as distinct targets of attention. For every user $i \in U_{\mathrm{mass}}$, we then construct an attention profile u_i: a vector of length k whose entries count the reposts that i directs toward each core block. Stacking these profiles yields, for every issue network, an $m \times k$ attention matrix A, where $m = |U_{\mathrm{mass}}|$, which forms the basis of all fragmentation analyses.

4.7 Quantifying fragmentation and its sorting principles

Null model for attentional patterns Previous research shows that elementary network properties, such as degree sequences, can impact observed network phenomena [39]. To ensure that our observed attention patterns are not due to differing popularity among distinct core voices, we introduce a controlled null model. We randomize the attention matrix A such that the total activity (row sum, or out-degree) of each mass user is exactly preserved, while maintaining the expected attention (column sum, or expected in-degree) for each core block. By doing so, we also control for varying activity levels, as low-activity users are intrinsically constrained in their potential for diversity compared to their high-activity peers.

This null constitutes the formal baseline for the single gaze. A literal single gaze does not imply that all observed footprints are identical, since users differ in how much they repost; rather, it implies that all users sample their endorsements from a *common* attention distribution over the core voices. Given each user's finite activity, observed footprints under such a model differ only through sampling noise, and this is precisely the scenario our randomization

reproduces: individual activity is preserved exactly, while attention is allocated in proportion to each core's overall popularity. The distribution of pairwise attention similarities under this null therefore captures the footprint diversity that a maximally unified system would still exhibit.

A common effect-size estimator All our fragmentation measures share the same non-parametric form. Given two populations of user pairs, $\mathcal{P}_1$ and $\mathcal{P}_2$, we compute the probability that a randomly drawn pair from $\mathcal{P}_1$ exhibits higher attentional similarity than a randomly drawn pair from $\mathcal{P}_2$:

$$\mathrm{CL}(\mathcal{P}_1, \mathcal{P}_2) = P\big(S_c(u_i, u_j) > S_c(u_x, u_y)\big), \quad (i,j) \in \mathcal{P}_1,\ (x,y) \in \mathcal{P}_2, \tag{3}$$

where S_c denotes the cosine similarity between attention profiles. This is the common language effect size (CLES), also known as the probability of superiority [63]: a non-parametric, outlier-robust, and interpretable measure. We linearly rescale it as $\delta = 2 \times (\mathrm{CL} - 0.5)$, which is mathematically equivalent to Cliff's Delta and ranges over $[-1, 1]$. A value of $\delta = 0$ indicates that the two similarity distributions are indistinguishable, while $\delta = 1$ indicates that every pair in $\mathcal{P}_1$ is more similar than any pair in $\mathcal{P}_2$.

Our three measures instantiate Eq. (3) with different pair populations. *Systemic fragmentation* Φ compares pairs drawn under the null model ($\mathcal{P}_1$) against observed pairs ($\mathcal{P}_2$): $\Phi = 0$ implies no evidence of fragmentation, whereas $\Phi = 1$ implies that every observed pair is less similar than any pair generated by the shared-gaze null, indicating a fragmented system. Negative values ($\Phi < 0$) would indicate hyper-convergence, i.e., observed pairs being systematically *more* similar than the shared-gaze baseline; we do not observe this in any of the twelve systems. To decompose how fragmentation is sorted, we partition the mass users either by their inferred attitudinal stance or by their structural block, and classify all user pairs as same-group or cross-group. *Attitudinal Cleavage* ϕ_A and *Factional Divergence* ϕ_F then compare same-group pairs ($\mathcal{P}_1$) against cross-group pairs ($\mathcal{P}_2$) under the respective partition, so that positive values indicate that the corresponding sorting principle structures the distribution of attention.

Monte Carlo estimation Exhaustive evaluation of Eq. (3) requires comparing all pairs across the two populations, each of which contains $O(m^2)$ pairs, resulting in $O(m^4)$ comparisons. We therefore approximate all three measures via Monte Carlo sampling, randomly drawing $N = 1{,}000{,}000$ pairs from each population to compute an unbiased estimator of the true effect size.

4.8 Null model for mainstream adherence

To assess the statistical significance of the inter-issue alignment and tensions captured by the *Mainstream Adherence* metric, we use a permutation-based null model. For any two issues, the analysis is restricted to the overlapping set of users active in both domains. We generate a null distribution by keeping user stances for the first issue fixed and randomly shuffling the stance labels for the second issue among the shared user base. This ensures the total counts of majority and minority stances for the second issue remain constant, which is necessary to account for the strong ideological homogeneity observed across the systems. By recalculating the metric across 10,000 permutations, we establish a baseline of expected adherence under a hypothesis of independence.

Acknowledgments

The calculations presented above were performed using computer resources within the Aalto University School of Science "Science-IT" project. We would also like to thank Dr. Dr. Corinna Coupette for their valuable assistance in proofreading and providing helpful comments on earlier drafts of this manuscript. We are grateful to Dr. Arttu Malkamäki for insightful discussions that helped shape the theoretical framework and terminology used in this work. This research was supported by grants from the Research Council of Finland (349366; 357743).

Ethics statement

This study analyzes exclusively public posts and interactions broadcast through Bluesky's public Firehose API, collected in accordance with the platform's terms of service. No private accounts, direct messages, or access-restricted content were collected. All analyses are conducted and reported at the aggregate level: no individual users are identified, and inferred user-level stance labels are used solely for aggregate statistical analysis and are not released.

References

[1] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

[2] Raymond S Nickerson. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, 2(2):175–220, 1998.

[3] Cass R. Sunstein. *Republic.com*. Princeton University Press, 2001.

[4] Cass Sunstein. The law of group polarization. *Journal of political philosophy*, 2002.

[5] W Lance Bennett and Shanto Iyengar. A new era of minimal effects? the changing foundations of political communication. *Journal of communication*, 58(4):707–731, 2008.

[6] Eli Pariser. *The filter bubble: What the Internet is hiding from you*. penguin UK, 2011.

[7] Richard Fletcher and Rasmus Kleis Nielsen. Are people incidentally exposed to news on social media? a comparative analysis. *New media & society*, 20(7):2450–2468, 2018.

[8] Pablo Barberá. How social media reduces mass political polarization. evidence from germany, spain, and the us. *Job Market Paper, New York University*, 46:1–46, 2014.

[9] Matthew Gentzkow and Jesse M Shapiro. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839, 2011.

[10] Jonathan Bright. Explaining the emergence of political fragmentation on social media: The role of ideology and extremism. *Journal of computer-mediated communication*, 23(1):17–33, 2018.

[11] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. Political polarization on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 5 (1), pages 89–96, 2011.

[12] Peter Dahlgren. The internet, public spheres, and political communication: Dispersion and deliberation. *Political Communication*, 22(2):147–162, 2005.

[13] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. Quantifying controversy on social media. *ACM Transactions on Social Computing*, 1(1):1–27, 2018.

[14] Jon Green, Jared Edgerton, Daniel Naftel, Kelsey Shoub, and Skyler J Cranmer. Elusive consensus: Polarization in elite communication on the covid-19 pandemic. *Science advances*, 6(28):eabc2717, 2020.

[15] Ali Salloum, Ted Hsuan Yun Chen, and Mikko Kivelä. Anatomy of elite and mass polarization in social networks. *Network Science*, 13:e18, 2025.

[16] Max Falkenberg, Fabiana Zollo, Walter Quattrociocchi, Jürgen Pfeffer, and Andrea Baronchelli. Patterns of partisan toxicity and engagement reveal the common structure of online political communication across countries. *Nature Communications*, 15(1):9560, 2024.

[17] Edoardo Di Martino, Alessandro Galeazzi, Michele Starnini, Walter Quattrociocchi, and Matteo Cinelli. Ideological fragmentation of the social media ecosystem: From echo chambers to echo platforms. *PNAS Nexus*, 4(9):pgaf262, 09 2025.

[18] Dorian Quelle, Frederic Denker, Prashant Garg, and Alexandre Bovet. Simple contagion drives population-scale platform migration, 2025.

[19] Dorian Quelle and Alexandre Bovet. Bluesky: Network topology, polarization, and algorithmic curation. *PloS one*, 20(2):e0318034, 2025.

[20] Yini Zhang, Josephine Lukito, Jiyoun Suk, and Ryan McGrady. Trump, twitter, and truth social: how trump used both mainstream and alt-tech social media to drive news media attention. *Journal of Information Technology & Politics*, 22(2):229–242, 2025.

[21] D.M. Anderson and M. Cornfield. *The Civic Web: Online Politics and Democratic Values*. Campaigning American style. Rowman & Littlefield, 2003.

[22] Natalie Fenton and John Downey. Counter public spheres and global modernity. *Javnost-The Public*, 10(1):15–32, 2003.

[23] Megan McArdle. The Bluesky bubble hurts liberals and their causes. *The Washington Post*, jun 2025. Accessed: 2026-06-24.

[24] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. The echo chamber effect on social media. *Proceedings of the national academy of sciences*, 118(9):e2023301118, 2021.

[25] Walter Quattrociocchi, Antonio Scala, and Cass R Sunstein. Echo chambers on facebook. *Available at SSRN 2795110*, 2016.

[26] Shanto Iyengar, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood. The origins and consequences of affective polarization in the united states. *Annual review of political science*, 22(1):129–146, 2019.

[27] Serge Moscovici and Marisa Zavalloni. The group as a polarizer of attitudes. *Journal of personality and social psychology*, 12(2):125, 1969.

[28] Jon C Rogowski and Joseph L Sutherland. How ideology fuels affective polarization. *Political behavior*, 38(2):485–508, 2016.

[29] Natalie Jomini Stroud. Polarization and partisan selective exposure. *Journal of communication*, 60(3):556–576, 2010.

[30] Stefan Schulz-Hardt, Dieter Frey, Carsten Lüthgens, and Serge Moscovici. Biased information search in group decision making. *Journal of personality and social psychology*, 78(4):655, 2000.

[31] Francoise Boucek. Rethinking factionalism: typologies, intra-party dynamics and three faces of factionalism. *Party politics*, 15(4):455–485, 2009.

[32] Brian Carey and Elizabeth Ventham. There is no fresh air: A problem with the concept of echo chambers. *Episteme*, page 1–19, 2025.

[33] Herbert A Simon. Bounded rationality. In *Utility and probability*, pages 15–18. Springer, 1990.

[34] Robert D Benford and David A Snow. Framing processes and social movements: An overview and assessment. *Annual review of sociology*, 26(2000):611–639, 2000.

[35] Jacquelien van Stekelenburg and Bert Klandermans. Dynamics of demand. In *A Social Psychology of Protest: Individuals in Action*, pages 61–97. Cambridge University Press, Cambridge, 2023.

[36] Robert Michels. Part vi, chapter 2. In Arthur Asa Berger, editor, *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. Routledge, 1999.

[37] Fabio Sani. Schism in groups: A social psychological account. *Social and Personality Psychology Compass*, 2(2):718–732, 2008.

[38] Peter Söderlund and Kimmo Grönlund. Can a change in the leadership of a populist radical right party be traced among voters? the case of the finns party. *Scandinavian Political Studies*, 47(1):23–46, 2024.

[39] Ali Salloum, Ted Hsuan Yun Chen, and Mikko Kivelä. Separating polarization from noise: comparison and normalization of structural polarization measures. *Proceedings of the ACM on human-computer interaction*, 6(CSCW1):1–33, 2022.

[40] Ryan J. Gallagher, Jean-Gabriel Young, and Brooke Foucault Welles. A clarified typology of core-periphery structure in networks. *Science Advances*, 7(12):eabc9800, 2021.

[41] Ted Hsuan Yun Chen, Ali Salloum, Antti Gronow, Tuomas Ylä-Anttila, and Mikko Kivelä. Polarization of climate politics results from partisan sorting: Evidence from finnish twittersphere. *Global Environmental Change*, 71:102348, 2021.

[42] Pedro Guerra, Wagner Meira Jr, Claire Cardie, and Robert Kleinberg. A measure of polarization on social media networks based on community boundaries. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7 (1), pages 215–224, 2013.

[43] Alfredo Jose Morales, Javier Borondo, Juan Carlos Losada, and Rosa M Benito. Measuring political polarization: Twitter shows the two sides of venezuela. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(3), 2015.

[44] Giulia Preti, Matteo Riondato, Aristides Gionis, and Gianmarco De Francisci Morales. Dsp: A statistically-principled structural polarization measure. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining*, pages 531–540, 2026.

[45] Martin Kleppmann, Paul Frazee, Jake Gold, Jay Graber, Daniel Holmgren, Devin Ivy, Jeromy Johnson, Bryan Newbold, and Jaz Volpert. Bluesky and the at protocol: Usable decentralized social media. In *Proceedings of the ACM Conext-2024 Workshop on the Decentralization of the Internet*, pages 1–7, 2024.

[46] Anne Oeldorf-Hirsch. X to bluesky platform migration: Governance and community. *AoIR Selected Papers of Internet Research*, 2025.

[47] Bluesky. 2025 Transparency Report Overview. Bluesky Blog, January 2026. Accessed: 2026-04-17.

[48] Mohsen Mosleh, Jennifer Allen, and David G Rand. Divergent patterns of engagement with partisan and low-quality news across seven social media platforms. *Proceedings of the National Academy of Sciences*, 122(44):e2425739122, 2025.

[49] Ingmar Weber, Venkata Rama Kiran Garimella, and Asmelash Teka. Political hashtag trends. In *European conference on information retrieval*, pages 857–860. Springer, 2013.

[50] Pew Research Center. Beyond red vs. blue: The political typology. Technical report, Pew Research Center, November 2021. Accessed: April 22, 2026.

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

[52] Harrison C White, Scott A Boorman, and Ronald L Breiger. Social structure from multiple networks. i. blockmodels of roles and positions. *American journal of sociology*, 81(4):730–780, 1976.

[53] Darko Hric, Kimmo Kaski, and Mikko Kivelä. Stochastic block model reveals maps of citation patterns and their evolution in time. *Journal of Informetrics*, 12(3):757–783, 2018.

[54] Elisabeth Noelle-Neumann. The spiral of silence a theory of public opinion. *Journal of communication*, 24(2):43–51, 1974.

[55] Bluesky. Firehose, 2025. Accessed: 2025-04-21.

[56] Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025.

[57] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. Umap: Uniform manifold approximation and projection. *Journal of Open Source Software*, 3(29):861, 2018.

[58] Ricardo JGB Campello, Davoud Moulavi, and Jörg Sander. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 160–172. Springer, 2013.

[59] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

[60] Geoffrey C Bowker and Susan Leigh Star. *Sorting things out: Classification and its consequences*. MIT press, 2000.

[61] Milagros Miceli, Martin Schuessler, and Tianling Yang. Between subjectivity and imposition: Power dynamics in data annotation for computer vision. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–25, 2020.

[62] Peter D Grünwald. *The minimum description length principle*. MIT press, 2007.

[63] Kenneth O McGraw and Seok P Wong. A common language effect size statistic. *Psychological bulletin*, 111(2):361, 1992.

[64] Letizia Iannucci. Sinitaivas live: Live streaming data collection from bluesky firehose, 2025.

[65] Aditya Kusupati, Gantavya Bhatt, Aniket Rege, Matthew Wallingford, Aditya Sinha, Vivek Ramanujan, William Howard-Snyder, Kaifeng Chen, Sham Kakade, Prateek Jain, et al. Matryoshka representation learning. *Advances in Neural Information Processing Systems*, 35:30233–30249, 2022.

[66] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.

A Data collection

A.1 Collection pipeline

Data was collected via a live stream from the *Bluesky Firehose*, an authenticated real-time stream of events encoded according to the AT Protocol. The collection pipeline decoded incoming CBOR events and stored them as raw Newline Delimited JSON (NDJSON) files. While the pipeline captures all available event types (posts, likes, follows) [64], this study exclusively filtered and utilized repost events for the construction of the interaction networks. The collection architecture ensured real-time data capture with automatic partition management for efficient querying.

A.2 Reliability and coverage

Data collection was interrupted only by two specific shutdowns: data loss on 16-01-2025 from 16:00 UTC until the end of the day, and a second outage starting at 11:00 UTC on 31-03-2025 that continued through the entire UTC days of 01-04-2025 and 02-04-2025. In total, we lost approximately 69 hours of events over the entire data collection period, resulting in an hourly uptime exceeding 99.2%.

A.3 Platform activity statistics

Table 5: Key platform activity statistics. Daily activity and total counts for *create* actions as recorded during our data collection, grouped by action type.

Action Type	Daily Average		Total Count
	Actions	Authors	Actions
Likes	29M	1.5M	9.3B
Posts	5.1M	0.8M	1.6B
Reposts	4.5M	0.5M	1.4B
Blocks	0.4M	0.1M	129M
Follows	5.7M	0.6M	1.8B
Sign-ups	—	46K	14M

During our data collection, we recorded the creation of 14M new profiles on the platform, an average of 46K sign-ups per day, which demonstrates the platform's increased popularity. User engagement is substantial: We logged over 1.6B new posts, over 1.4B reposts, and over 9.3B likes. All statistics are listed in Table 5.

B Topic inference pipeline

The pipeline identifies latent thematic structures by combining high-dimensional semantic embeddings with density-based clustering, removing the need for predefined keywords or any prior knowledge.

Embedding and dimensionality reduction We first generated text embeddings using *Qwen/Qwen3-Embedding-8B* [56]. To balance computational efficiency with semantic depth, we utilized *Matryoshka Representation Learning* [65] to truncate the original 4096-dimensional outputs to 248 dimensions. These vectors were subsequently projected into a 12-dimensional manifold via UMAP [57] to prepare the data for structural density estimation.

Clustering and optimization Dense regions within this manifold were identified using HDBSCAN [58]. Because the resulting topology is sensitive to input parameters, we executed a grid search across 200 distinct configurations. The final set of parameters was selected by maximizing cluster granularity (obtaining a Silhouette score ≈ 0.48) among the top-performing candidates (see Figure 6).

Interpretation and labeling To transform these clusters into interpretable topics, we used *Qwen/Qwen3-30B-A3B-Instruct-2507* served via vLLM [66]. For each cluster, we sampled 100 representative posts and generated descriptive labels and thematic summaries using the prompt in Listing B. A full breakdown of the optimal hyperparameters and LLM generation settings is provided in Table 6.

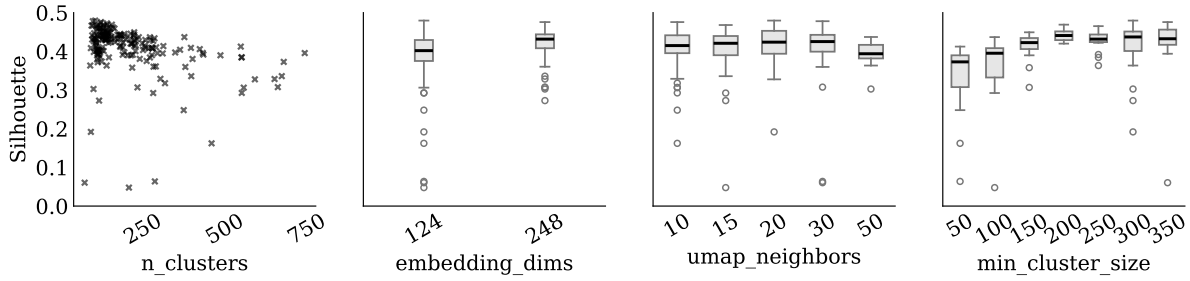


Figure 6: Grid search results. Silhouette scores for 200 configurations across UMAP (components 10–20, neighbors 10–50) and HDBSCAN (min_cluster 50–350, ϵ 0–0.15) parameters.

Table 6: Optimal pipeline hyperparameters. Selected configurations for dimensionality reduction (UMAP), density-based clustering (HDBSCAN), and label generation (SUMMARIZE).

Step	Parameter	Value
UMAP	`n_components`	12
	`n_neighbors`	10
	`min_dist`	0.05
	`metric`	cosine
HDBSCAN	`min_cluster_size`	350
	`min_samples`	100
	`cluster_selection_epsilon`	0.15
	`metric`	euclidean
	`selection_method`	eom
SUMMARIZE	`temperature`	0.3
	`max_output`	128
	`context`	100 sampled posts
	`model`	Qwen3-30B-A3B-Instruct-2507

System Prompt: Cluster Summarization

Role: You are given a set of short social media posts that belong to the same cluster. Your task is to produce:

1. Title A short, descriptive title (max 10 words) capturing the main theme.

2. Description A short description (max 2 sentences) summarizing the content.

Instructions

- If posts touch on many topics, focus on the most prevalent theme.
- Keep wording neutral and concise.

Posts:
`[cluster_posts]`

C Grouping of microtopics

The previously described pipeline for preliminary topic exploration found 122 microtopics from our dataset, and they serve as the foundational basis for constructing our final high-level issue categories. As many emerging microtopics appear to be thematically close to each other, we first organize them into six higher-level political categories and seven for non-political topics to obtain a better overview on prevalent core themes. Full descriptive summaries and TF-IDF-based keywords for these microtopics are provided in the Supplementary Information.

I. POLITICS

US Politics: Trump-Centric & Administration
8: Project 2025: The Blueprint for Authoritarianism · **17:** Trump's Tariffs and the Farmer Bailout Crisis · **20:** Trump's Greenland Obsession Sparks Global Outrage · **32:** Trump's Attack on Public Education · **33:** Pete Hegseth Scandal: Security & Incompetence at the Pentagon · **34:** Charlie Kirk's Murder and Right-Wing Backlash · **47:** Viral Outrage Over JD Vance's Controversial Persona · **73:** Trump's Jan. 6 Pardons and Political Fallout · **78:** Trump's Military Deployments and Martial Law Fears · **79:** ICE Raids and Mass Deportations Under Trump · **81:** Trump's Tariff Wars Spark Global Economic Fears · **91:** Republican Leaders Accused of Pedophilia · **92:** Demand for Release of Epstein Files · **93:** Trump's Decline and Cultural Backlash · **96:** Social Security Under Attack · **97:** 2024 Election Rigging Allegations · **101:** Elon Musk's Government Overhaul: Chaos or Conspiracy? · **106:** Call for Trump's Third Impeachment

US Politics: Parties, Institutions & Ideology
83: MAGA Critique and Cultural Backlash · **85:** GOP-Controlled Shutdown Crisis · **89:** Fascism Is Here: Resistance Is Urgent · **90:** Nazi Parallels in Modern Politics · **100:** Medicaid Under Attack: GOP Cuts vs. Public Health · **105:** SCOTUS Under Fire: Crisis of Constitutional Authority · **108:** Calls for Schumer's Resignation Amid Leadership Crisis · **109:** Democratic Party Crisis and Calls for Radical Change

International Politics
4: Pope Francis: Legacy and the New Papacy · **61:** Ukraine's Struggle for Sovereignty and Western Support · **86:** UK Political Realignment and Labour's Challenges · **110:** Canada Stands Firm Against Threats · **111:** Mark Carney's Rise and the Fall of Poilievre · **112:** Genocide in Gaza and Global Complicity

Social Issues & Identity
5: Celebrating Black and Women's Histories · **11:** Reproductive Rights Under Siege · **14:** DEI Backlash: Civil Rights Under Attack · **19:** Pride Beyond June: Resistance and Joy · **26:** Autism, Neurodiversity, and the Fight for Recognition · **36:** Critique of Christian Nationalism and Hypocrisy · **39:** Crisis of American Identity and National Shame · **87:** Trans Rights, Misogyny, and the War on Gender Identity · **88:** White Privilege and Accountability in Racial Discourse

Public Policy, Activism & Environment
6: Aviation Safety Crisis Amid Political and Corporate Chaos · **22:** Mass Consumer Boycotts Against Corporate Power · **31:** Nationwide 'No Kings' Protests Against Authoritarianism · **58:** FEMA & NOAA Cuts Amid Deadly Floods · **59:** RFK Jr. and the Crisis of Public Health Trust · **66:** Global Wildfires and Climate Crisis · **67:** Renewables Revolution: Solar, Storage, and the Grid Shift · **71:** Housing Crisis: Supply, Power, and Policy Failures · **72:** Global Push for Safe, Active Transportation · **99:** Tax the Rich: A Call for Fairness · **107:** Flood the Switchboard: Call Your Senators Now

Corporate, Tech & Media Politics
37: Media Complicity in the Erosion of Truth · **45:** AI Anxiety and the Crisis of Authenticity · **102:** Tesla Boycott and Protest Movement

II. NON-POLITICS

Entertainment & Media
TV, Film & Celebrities: **1:** BB27: Chaos, Alliances, and Vorgan Hours · **7:** Drama, Betrayal, and Family Feuds in #BeyondTheGates · **65:** Obscure Films & Cult Cinema Passion
Music & Podcasts: **16:** BTS Reunion Hype & Global K-Pop Momentum · **43:** Live Music & Concert Announcements · **50:** Music Nostalgia & Album Reflections · **64:** Podcast Launches & Updates 2025
Sports: **9:** AEW Dynamite & Collision Highlights · **51:** Sports Highlights & Fan Reactions · **52:** F1 Drama, Cycling Triumphs, and Motorsport Highlights

Gaming, Fandom & Hobbies
Video Games: **74:** Virtual Photography in Video Games · **80:** FFXIV G-Pose Community & Glamour Culture · **114:** Indie Game Development Journey · **115:** Live Gaming & VTuber Streams
Analog Hobbies: **24:** Car Culture & Classic Rides · **30:** Marvel & DC Comics Culture · **44:** Transformers & Anime Nostalgia · **75:** Warhammer Miniature Painting & Kitbashing Journey · **84:** Handmade Crafts & Fiber Arts Journey
Fan Creations: **29:** Deltarune Fan Art & Ship Culture · **40:** Hazbin Hotel Fan Creations & Season 2 Hype · **69:** Indie Comics & Webcomic Creators · **70:** Fanfic Updates & Recs Galore · **95:** ArtFight 2025: Creative Attacks & Revenants · **98:** Fursuit Adventures & Community Fun · **103:** Furry & OC Art Showcase · **104:** Pokémon Art & Fan Creativity

Online Content & Platforms
Video Platforms: **0:** Compilation of Viral and Disturbing Online Videos · **2:** TikTok Culture & Viral Moments · **116:** Recommended YouTube Videos · **117:** YouTube Shorts Compilation
Writing & Lit: **12:** Substack Highlights & Newsletter Roundup · **18:** NYT Gift Articles & News Highlights · **25:** Echoes of Longing and Nature (Poetry) · **53:** Celebrating Queer & Indie Voices in Books
Memes & Culture: **3:** The Urgent Need for Alt Text on Bluesky · **13:** Meme Culture & Community Sharing · **49:** Toxic Behavior and Denial in Online Spaces · **57:** Bluesky Community Dynamics and Migration

Science & Technology (Academic)
28: AI and Machine Learning Research Highlights · **46:** Open Source, Linux, and Dev Tools Trends · **60:** Science, Research, and Academic Opportunities

Personal Life & Lifestyle
Daily Life: **10:** Happy New Year 2025 Wishes · **15:** Valentine's Day Celebrations & Self-Love · **21:** Birthday Celebrations & Wishes · **41:** The Eternal Battle with Sleep · **42:** Hump Day Vibes & Daily Greetings · **56:** Gratitude and Appreciation Exchange · **77:** Hair, Identity, and the Ritual of Transformation
Food, Hobbies & Pets: **38:** Cities, Pride, and Wanderlust · **62:** Food Obsessions & Kitchen Chaos · **63:** Gym Life: Struggles & Triumphs · **120:** Heartwarming & Hilarious Dog Moments · **121:** Caturday Chaos & Cat Shenanigans
Nature: **54:** Severe Weather & Storm Alerts Worldwide · **68:** Weather Whiplash · **118:** Birdwatching & Wildlife Photography · **119:** Nature, Landscapes, and the Art of Photography

Crowdfunding & Community Support
35: Art Commissions Open – Limited Slots · **48:** Creator Support & Patreon Rewards · **55:** Urgent Animal Rescues & Medical Appeals · **76:** Urgent Mutual Aid Requests · **94:** Indie Creative Kickstarters in Final Stretch · **113:** Gaza Families Beg for Survival

Explicit & Niche Content
23: Fine Art Nude Photography & Modeling · **27:** ABDL Lifestyle & Diaper Culture · **82:** NSFW Content & Fetish Expressions

Based on both the fine-grained microtopics and their aggregation into higher-level categories, and considering their relative prevalence, we then selected a final set of issues for in-depth analysis. We believe that these twelve issues are a good representation of the range of political topics that emerged organically from the corpus.

Issue	Description	Related Clusters
Trump Administration	Policies and actions of US President Donald Trump.	4, 8, 17, 20, 32, 33, 34, 47, 73, 78, 79, 81, 91, 92, 93, 96, 97, 106
Elon Musk	Political and regulatory issues concerning Elon Musk.	101, 102
Republican Party	Views on the US Republican Party.	83, 85, 89, 90, 91, 100, 105
Democratic Party	Views on the US Democratic Party.	83, 85, 89, 90, 91, 100, 105, 108, 109
Russia–Ukraine Conflict	Ongoing conflict between Russia and Ukraine.	61
Israel–Palestine Conflict	Ongoing conflict between Israel and Palestine.	112, 113
Reproductive Rights	Legal and ethical status of abortion.	11
Gender & Sexuality Rights	Rights and treatment of the LGBTQ+ community.	19, 87
Equity & Representation	Diversity, Equity, and Inclusion (DEI) programs.	14, 88
Climate & Environment	Climate change and environmental regulations.	58, 66, 67, 71, 72
Artificial Intelligence Policy	AI regulation and societal impact.	28, 45
Technology Platforms Policy	Regulation of social media and tech platforms.	3, 37, 49, 57, 101, 102

Table 7: Salient political issues with short descriptions and related clusters

D Codebook as a prompt for issue annotation

System Prompt: Social Media Post Classifier

Role: You are a social media post classifier. Your task is to read a post and assign it to exactly one category from the list below.

Categories

Trump Administration Posts about the Trump presidency or its officials, policies, actions, investigations, or consequences. Includes references to the administration's decisions, executive orders, legislative agenda, agency leadership, or comparisons explicitly tied to that period. Use this category if the post is clearly about Trump, whether about his government, his campaigns, or his role in U.S. politics.

Elon Musk Posts about Elon Musk in a U.S. governmental or political context. Includes congressional hearings, agency investigations, regulatory disputes, policy debates, or government contracts. Use this category if the post discusses Elon Musk in connection to U.S. government decisions, oversight, or policy debates.

Republican Party Posts expressing supportive or critical views about the US Republican Party's general platform, leadership, or political brand. Use this category if the post is about the US Republican Party specifically.

Democratic Party Posts expressing supportive or critical views about the US Democratic Party's general platform, leadership, or political brand. Use this category if the post is about the US Democratic Party specifically.

Russia–Ukraine Conflict Posts about the Russia–Ukraine conflict. Covers references to Russia, Ukraine, Kyiv, Kremlin, Zelensky, Putin, Donbas, Crimea, NATO, invasion, territorial disputes, and European/US military aid. Includes hostilities, ceasefires, casualties, diplomacy, sanctions, humanitarian aid, aid packages, or official government/organization statements.

Israel–Palestine Conflict Posts about the Israel–Palestine Conflict. Covers references to Israel, Palestine, Gaza, West Bank, Hamas, Netanyahu, IDF, settlers, occupation, UNRWA, and related peace talks. Includes hostilities, ceasefires, casualties, diplomacy, sanctions, humanitarian aid, aid packages, or official government/organization statements.

Reproductive Rights Posts about laws, court cases, news, policies, or government actions concerning abortion rights (pro-choice/pro-life) and other reproductive health legislation.

Gender and Sexuality Rights Posts about laws, court cases, news, policies, or government actions concerning LGBTQ+ people (including transgender rights). Includes anti-discrimination protections, education/healthcare access, marriage/adoption, bathroom/sports policies.

Equity and Representation Policy (DEI) Posts about organizational or governmental DEI programs (Diversity, Equity, and Inclusion). Includes debates, and policies (hiring/promotion practices, affirmative action, training, audits, or legal challenges) in institutional contexts.

Artificial Intelligence Policy Posts about AI regulation, governance, or policy debates tied to privacy or intellectual property. Includes rulemaking, standards, compliance, lawsuits, training-data/copyright disputes, and data-protection enforcement.

Technology Platforms Policy Posts about government actions, laws, court rulings, or hearings regarding social media and other tech platforms. Includes bans/restrictions, content moderation mandates, data security/privacy rules, or antitrust enforcement.

Climate and Environmental Policy Posts about climate change, environmental regulations, or renewable energy policies.

Other Political Posts that are political or public-policy related but do not fit any specific category above.

Non-Political Covers posts not related to politics or policy. Includes entertainment, sports, personal life, culture, memes, gaming, marketing, lifestyle, and other non-political content.

Classification rules

- Assign the post to exactly one category. Prefer the most specific applicable category.
- If a post clearly relates to multiple specific categories, use the one that represents the primary focus in the text.
- Do not infer unstated facts. Use only the content of the post (hashtags/mentions/emojis count as content).
- If uncertainty remains, use *Other Political* or *Non-Political*.

E Evaluation of issue detection

Table 8: Classification performance metrics by issue. The table shows precision, recall, and F1-score for each category in the model's predictions (NCC threshold = 0.62, Coverage = 90.1%).

Issue	Precision	Recall	F1-Score
Trump Administration	0.889	0.947	0.917
Elon Musk	1.000	1.000	1.000
Republican Party	0.909	0.909	0.909
Democratic Party	1.000	0.833	0.909
Russia-Ukraine Conflict	0.929	0.929	0.929
Israel-Palestine Conflict	1.000	1.000	1.000
Reproductive Rights	1.000	1.000	1.000
Gender and Sexuality Rights	0.818	0.900	0.857
Equity and Representation Policy	1.000	1.000	1.000
Artificial Intelligence Policy	0.714	1.000	0.833
Technology Platforms Policy	0.667	1.000	0.800
Climate and Environmental Policy	1.000	1.000	1.000
Other Political	0.944	0.580	0.718
Non-Political	0.952	0.991	0.971
Accuracy			0.941
Macro Avg	0.916	0.935	0.917
Weighted Avg	0.943	0.941	0.937

F Stance distributions with different thresholds

Table 9: Stance distribution across political issues. The table shows stance labels and user percentages per issue. Percentages are calculated for users with at least **two** posts or reposts on the issue.

	Majority Stance		Minority Stance		Neutral / Other
Issue Network	*Label*	*%*	*Label*	*%*	*%*
Trump Administration	Critical	97.68	Supportive	0.53	1.79
Elon Musk	Critical	98.52	Supportive	0.15	1.32
Republican Party	Critical	98.69	Supportive	0.23	1.08
Democratic Party	Supportive	48.01	Critical	47.68	4.32
Russia-Ukraine	Pro-Ukraine	91.72	Pro-Russia	1.78	6.50
Israel-Palestine	Pro-Palestine	95.31	Pro-Israel	2.26	2.43
Reproductive Rights	Pro-Choice	93.14	Pro-Life	1.40	5.46
Gender and Sexuality	Supportive	94.47	Critical	4.49	1.04
Equity and Representation	Supportive	83.94	Critical	13.82	2.24
Climate and Environmental	Activist	92.37	Dismissive	1.29	6.34
Artificial Intelligence	Critical	89.99	Supportive	3.54	6.47
Technology Platforms	Critical	96.92	Supportive	1.24	1.84

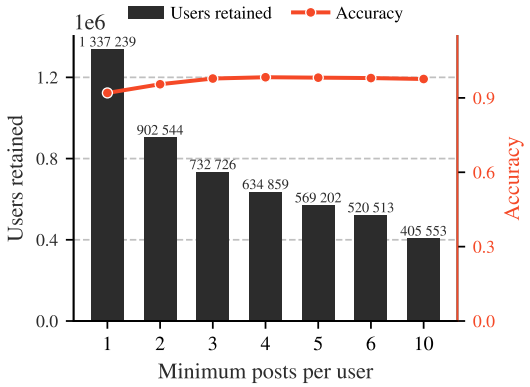
Table 10: Stance distribution across political issues. The table shows stance labels and user percentages per issue. Percentages are calculated for users with at least **ten** posts or reposts on the issue.

	Majority Stance		Minority Stance		Neutral / Other
Issue Network	*Label*	*%*	*Label*	*%*	*%*
`Trump Administration`	Critical	99.68	Supportive	0.16	0.16
`Elon Musk`	Critical	99.95	Supportive	0.03	0.02
`Republican Party`	Critical	99.83	Supportive	0.06	0.11
`Democratic Party`	Supportive	54.34	Critical	44.87	0.79
`Russia-Ukraine`	Pro-Ukraine	98.26	Pro-Russia	0.60	1.14
`Israel-Palestine`	Pro-Palestine	97.92	Pro-Israel	1.67	0.41
`Reproductive Rights`	Pro-Choice	99.71	Pro-Life	0.11	0.18
`Gender and Sexuality`	Supportive	98.95	Critical	0.99	0.06
`Equity and Representation`	Supportive	93.45	Critical	6.39	0.15
`Climate and Environmental`	Activist	99.17	Dismissive	0.16	0.67
`Artificial Intelligence`	Critical	97.28	Supportive	1.43	1.29
`Technology Platforms`	Critical	99.84	Supportive	0.12	0.04

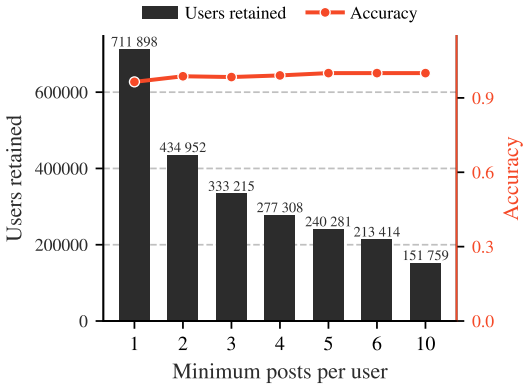
G Evaluation of stance detection

To evaluate the user-level stance classification, we manually annotated a random sample of 200 users per issue. For each sampled user, the annotators reviewed the same collection of issue-relevant posts and reposts that was provided to the classifier (see Methods 4.4) and independently assigned a stance label; accuracy is the share of users for whom the classifier's label matches this manual annotation.

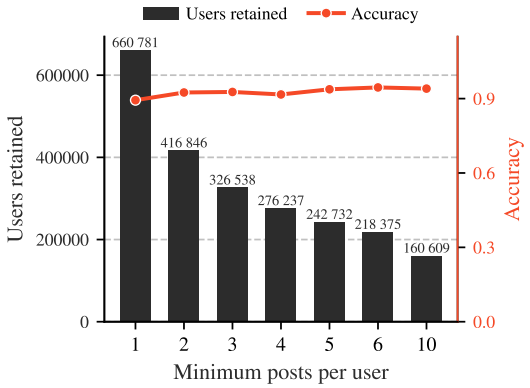
Figure 7 reports, for each of the twelve issues, classification accuracy as a function of the minimum number of issue-relevant posts or reposts required for a user to be included, together with the number of users retained at each threshold. Accuracy improves with user activity, as more content provides a stronger signal of a user's leaning, while stricter thresholds reduce coverage. The threshold of at least five posts or reposts used in the main analysis (Table 2) balances label reliability against population coverage; Appendix F reports the corresponding stance distributions at thresholds of two and ten.



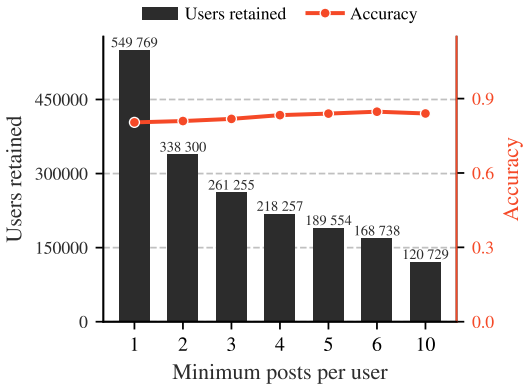
(A) TRUMP



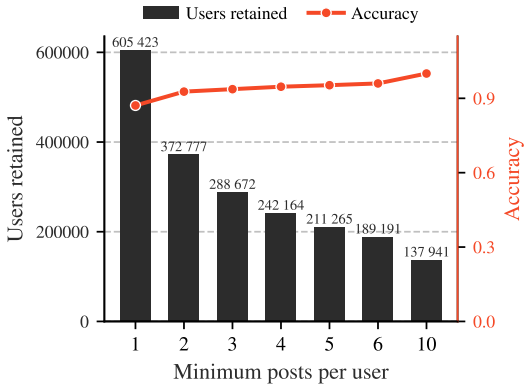
(B) MUSK



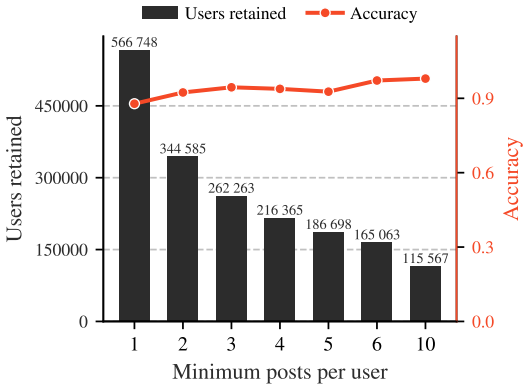
(C) GOP



(D) DEM

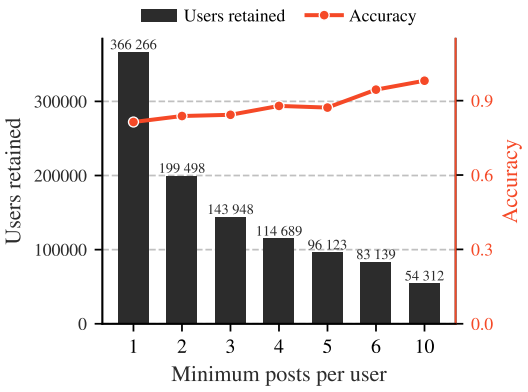


(E) RUS-UKR

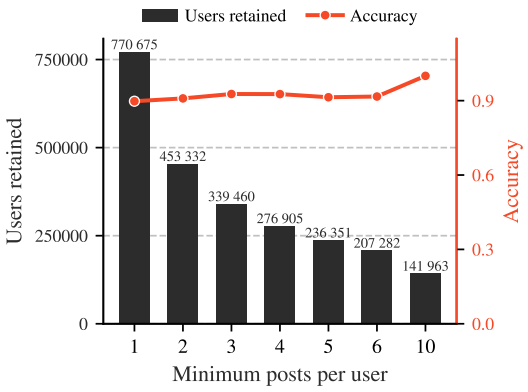


(F) ISR-PAL

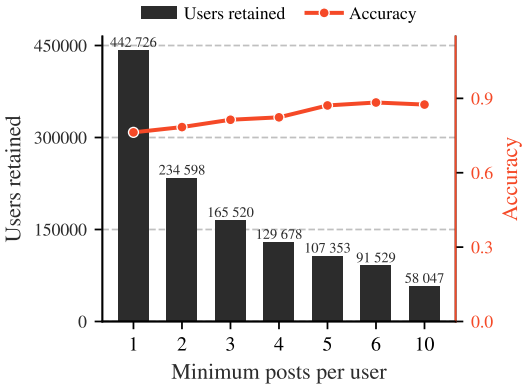
Figure 7: Stance classifier performance across 12 issues (Part 1/2).



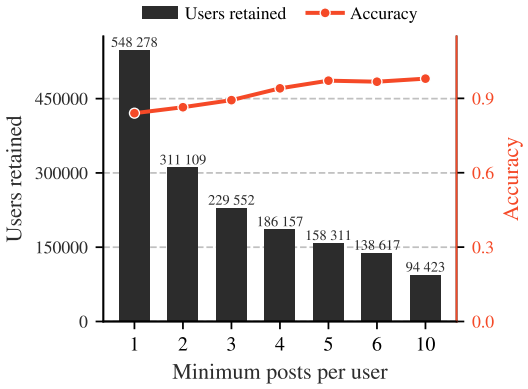
(G) REPRO



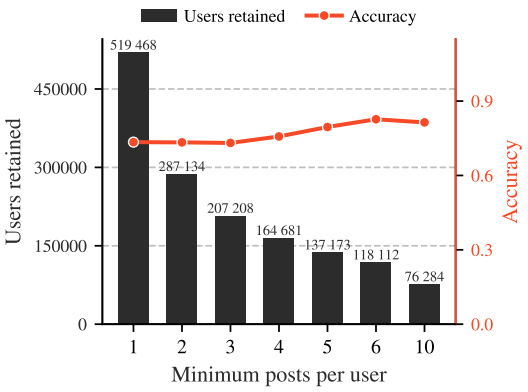
(H) GSR



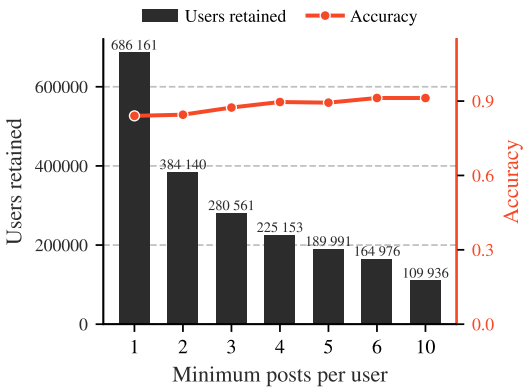
(I) ERP



(J) CLIMATE



(K) AI



(L) TECH

Figure 7: Stance classifier performance across 12 issues (Part 2/2).

H Connectivity patterns at distinct scales

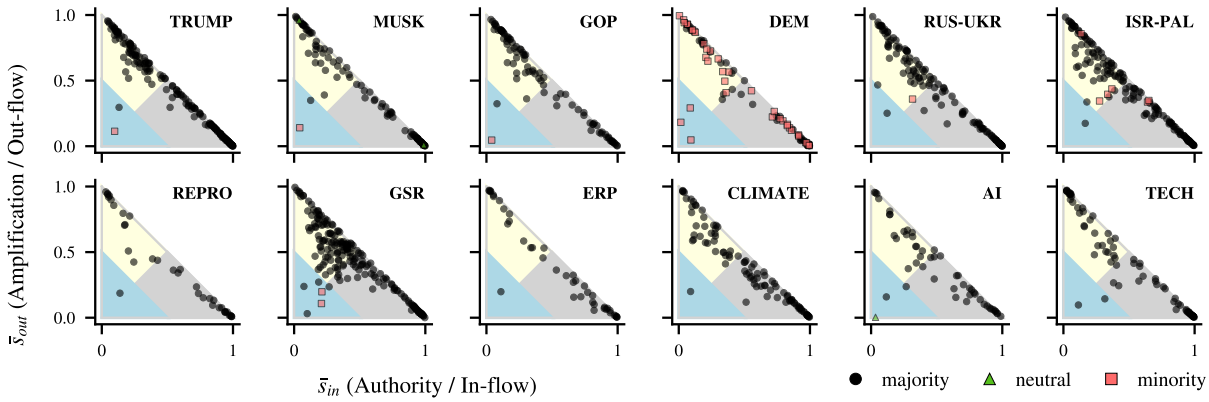


Figure 8: Structural connectivity patterns at level 1 of hSBM hierarchy across twelve networks. This figure displays the properties of blocks identified by the nested stochastic block model, mapped by their authority in-flow ($\bar{s}_{in}$) and amplification out-flow ($\bar{s}_{out}$). Because the sum of these flows is ≤ 1, all blocks are contained within the triangle bounded by the diagonal.

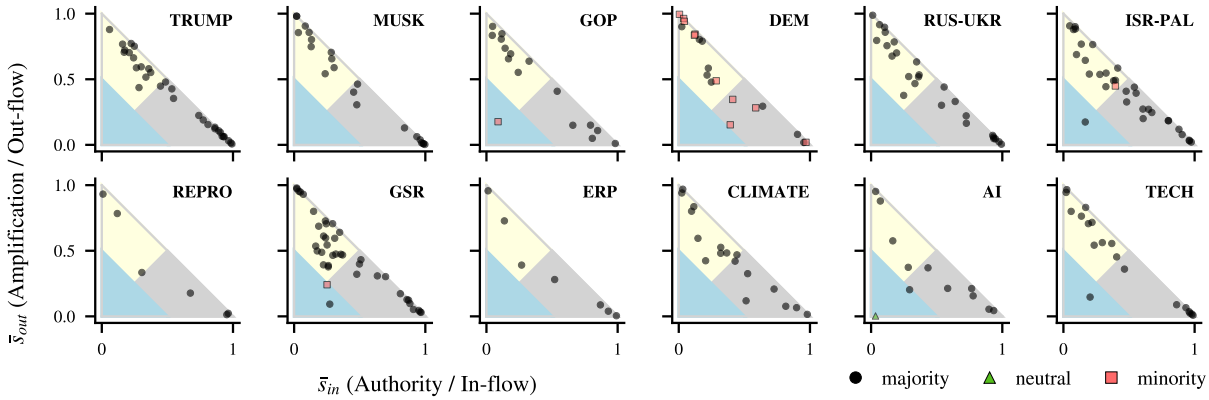


Figure 9: Structural connectivity patterns at level 3 of hSBM hierarchy across twelve networks. This figure displays the properties of blocks identified by the nested stochastic block model, mapped by their authority in-flow ($\bar{s}_{in}$) and amplification out-flow ($\bar{s}_{out}$). Because the sum of these flows is ≤ 1, all blocks are contained within the triangle bounded by the diagonal.

I Prevalence of US-based liberal-progressive and anti-conservative sentiment

The microtopics identified in the corpus show a high degree of congruence with the policy priorities and values of the *Progressive Left*, *Outsider Left*, and *Establishment Liberals* as defined by the Pew Research Center [50]. By mapping these findings to specific typology statistics, we can contextualize the observed discourse within the broader U.S. political landscape. Our manual review of the individual posts associated with these microtopics further clarifies the attitudes and rhetorical tone prevalent in the user base.

Trump-Centric Discourse and Institutional Critique (e.g., microtopics 8, 73, 79, 83, 89) Topics such as *Project 2025*, *Jan. 6 Pardons*, and *MAGA Critique* map directly to the Progressive Left, Establishment Liberals, and Outsider Left, who hold the most negative views of the U.S. president Trump and the Republican Party. On a 'feeling thermometer' (0–100), the Progressive Left gives Trump a nearly maximum cold rating of 2, while the Establishment Liberals gives him a 7, and the Outsider Left a 14. The respective ratings towards the Republican Party are, 10, 20 and 27. Regarding immigration policies, the actions of ICE clash with the left's views; 86% of the Progressive Left say that immigrants living in the U.S. illegally make their communities better. Furthermore, 63% of the Progressive Left and 54% of the Outsider Left believe the U.S. should actively increase the number of legal immigrants it admits.

Corporate Power and Economic Policy (e.g., microtopics 22, 99, 100, 102) Topics dealing with corporate boycotts, progressive taxation, and the social safety net (Medicaid) are issues where the Democratic coalition is in near-total agreement, though the Progressive Left is the most extreme. There is widespread agreement across all Democratic-oriented groups that tax rates on large businesses, corporations, and households making over 400,000 dollars should be raised. However, the Progressive Left is the most unified on this issue, with an overwhelming 98% supporting higher tax rates on large businesses and corporations. Almost all members of the Democratic coalition agree that the U.S. economic system 'unfairly favors powerful interests', including 97% of the Progressive Left and 94% of the Outsider Left. However, the Progressive Left is distinctly more critical of extreme wealth: 73% of the Progressive Left believe the existence of billionaires is a 'bad thing for the country', compared to 55% of the Outsider Left, and 39% of Establishment Liberals.

Social Identity and Systemic Justice (e.g., microtopics 14, 19, 36, 87, 88) The Progressive Left is supportive of gender identity issues, with 88% saying that greater social acceptance of transgender people is 'good for the country'. Moreover, 96% of the Progressive Left say 'a lot more' needs to be done to 'ensure equal rights', and 88% see a great deal of discrimination against Black people. As a point of comparison, only 4% of *Faith and Flag Conservatives* agree on this. Most notably, 71% of the Progressive Left believe that most U.S. laws and institutions 'need to be completely rebuilt' because they are 'fundamentally biased against' certain racial and ethnic groups.

Climate Crisis and Energy Transition (e.g., microtopics 66, 67) Groups aligned with the Democratic Party broadly agree that climate change is a significant national issue that justifies the cost of tougher environmental laws and a priority on alternative energy sources such as wind and solar. However, differences in intensity become clear when examining specific policies. For instance, 78% of the Progressive Left support a total phase-out of oil, coal, and natural gas. This position is supported by a smaller majority of the Outsider Left at 59% and approximately half of Establishment Liberals at 51%.

The alignment between the emerged microtopics and the Pew Research Center's typology [50] suggests that the platform's discourse is weighted toward the specific policy preferences and social perspectives of the U.S. political left. The consistency of this alignment across diverse domains, including institutional trust, economic structure, social justice, and environmental policy, shows that the observed discourse reflects a worldview most closely associated with the Progressive Left and its adjacent cohorts, specifically Establishment Liberals and Outsider Left.

J Fragmentation results

Table 11: The table reports the mean fragmentation indices (Cliff's Delta values) and 95% confidence intervals for global Systemic Fragmentation (Φ), Attitudinal Cleavage (ϕ_A), and Factional Divergence (ϕ_F) across each of the twelve constructed networks, sorted in descending order by Systemic Fragmentation.

Issue Network	Systemic Fragmentation *[95% CI]*	Attitudinal *[95% CI]*	Factional *[95% CI]*
`Trump Administration`	**0.273** *[0.272, 0.275]*	**0.520** *[0.519, 0.522]*	**0.185** *[0.184, 0.187]*
`Elon Musk`	**0.225** *[0.223, 0.227]*	**0.324** *[0.323, 0.325]*	**0.228** *[0.226, 0.229]*
`Russia-Ukraine Conflict`	**0.223** *[0.221, 0.225]*	**0.301** *[0.300, 0.303]*	**0.253** *[0.252, 0.255]*
`Republican Party`	**0.216** *[0.215, 0.218]*	**0.554** *[0.552, 0.555]*	**0.251** *[0.249, 0.252]*
`Technology Platforms`	**0.203** *[0.201, 0.204]*	**0.263** *[0.261, 0.264]*	**0.251** *[0.250, 0.252]*
`Gender and Sexuality Rights`	**0.190** *[0.188, 0.192]*	**0.068** *[0.067, 0.070]*	**0.398** *[0.397, 0.400]*
`Artificial Intelligence`	**0.187** *[0.185, 0.189]*	**0.336** *[0.335, 0.338]*	**0.429** *[0.427, 0.430]*
`Democratic Party`	**0.177** *[0.175, 0.179]*	**0.200** *[0.198, 0.201]*	**0.282** *[0.281, 0.284]*
`Equity and Representation`	**0.171** *[0.168, 0.173]*	**0.103** *[0.102, 0.105]*	**0.218** *[0.217, 0.219]*
`Israel-Palestine Conflict`	**0.167** *[0.165, 0.169]*	**0.143** *[0.142, 0.144]*	**0.365** *[0.364, 0.366]*
`Climate and Environment`	**0.137** *[0.135, 0.139]*	**0.042** *[0.041, 0.044]*	**0.344** *[0.342, 0.345]*
`Reproductive Rights`	**0.121** *[0.118, 0.123]*	**0.152** *[0.150, 0.153]*	**0.294** *[0.293, 0.295]*

K Attentional similarity distributions

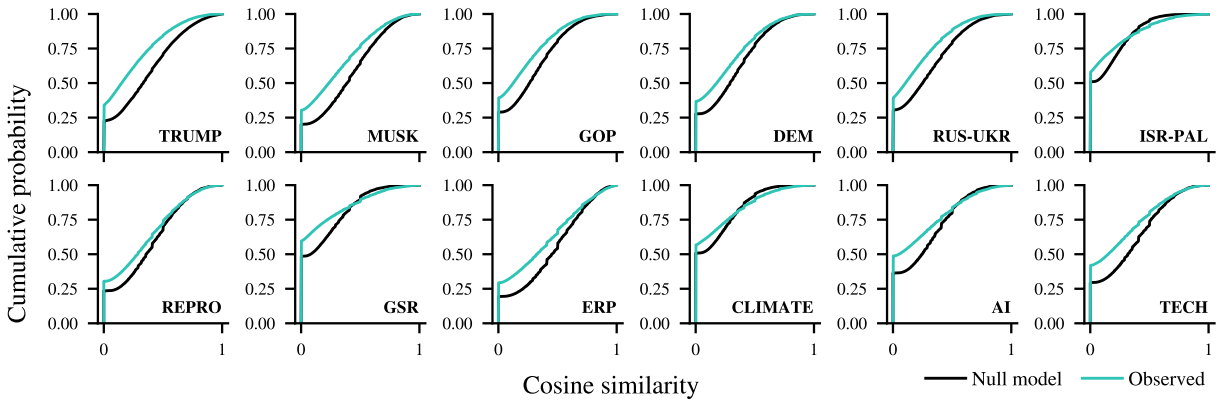


Figure 10: Cumulative attentional divergence from thel null model.

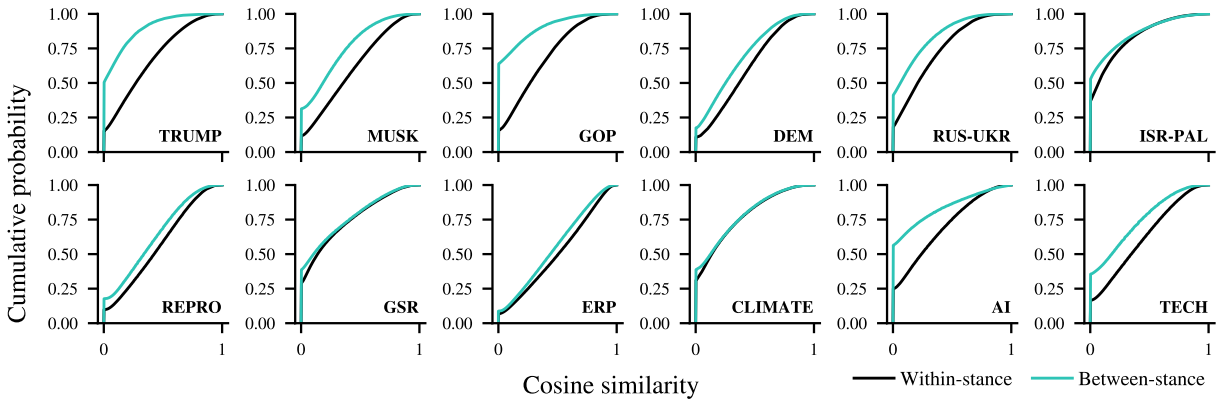


Figure 11: Attentional similarity distributions by ideological stance. Histograms comparing users who share the same stance (teal) against those with opposing stances (gray).

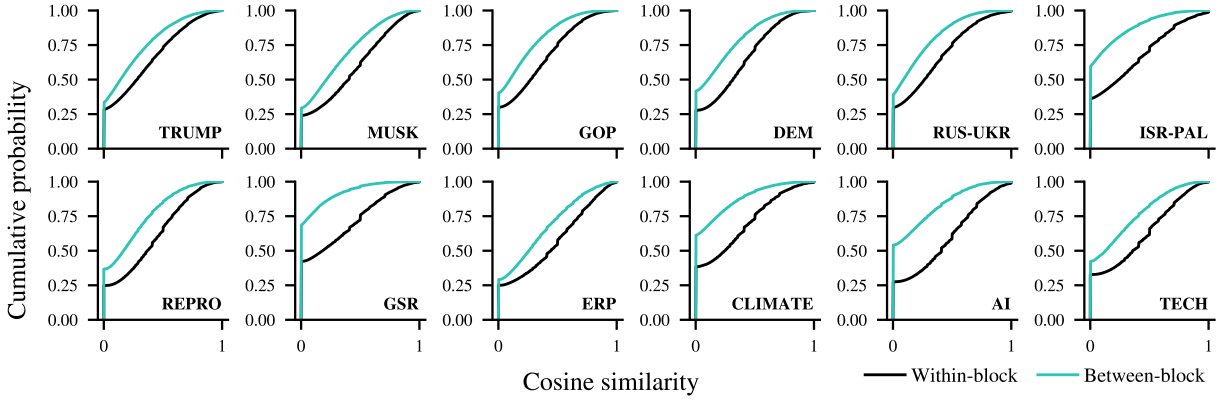


Figure 12: Attentional similarity distributions by hSBM block. Histograms of pairwise attention similarity comparing users within the same structural block (teal) against those in different blocks (gray).

L Minority alignment

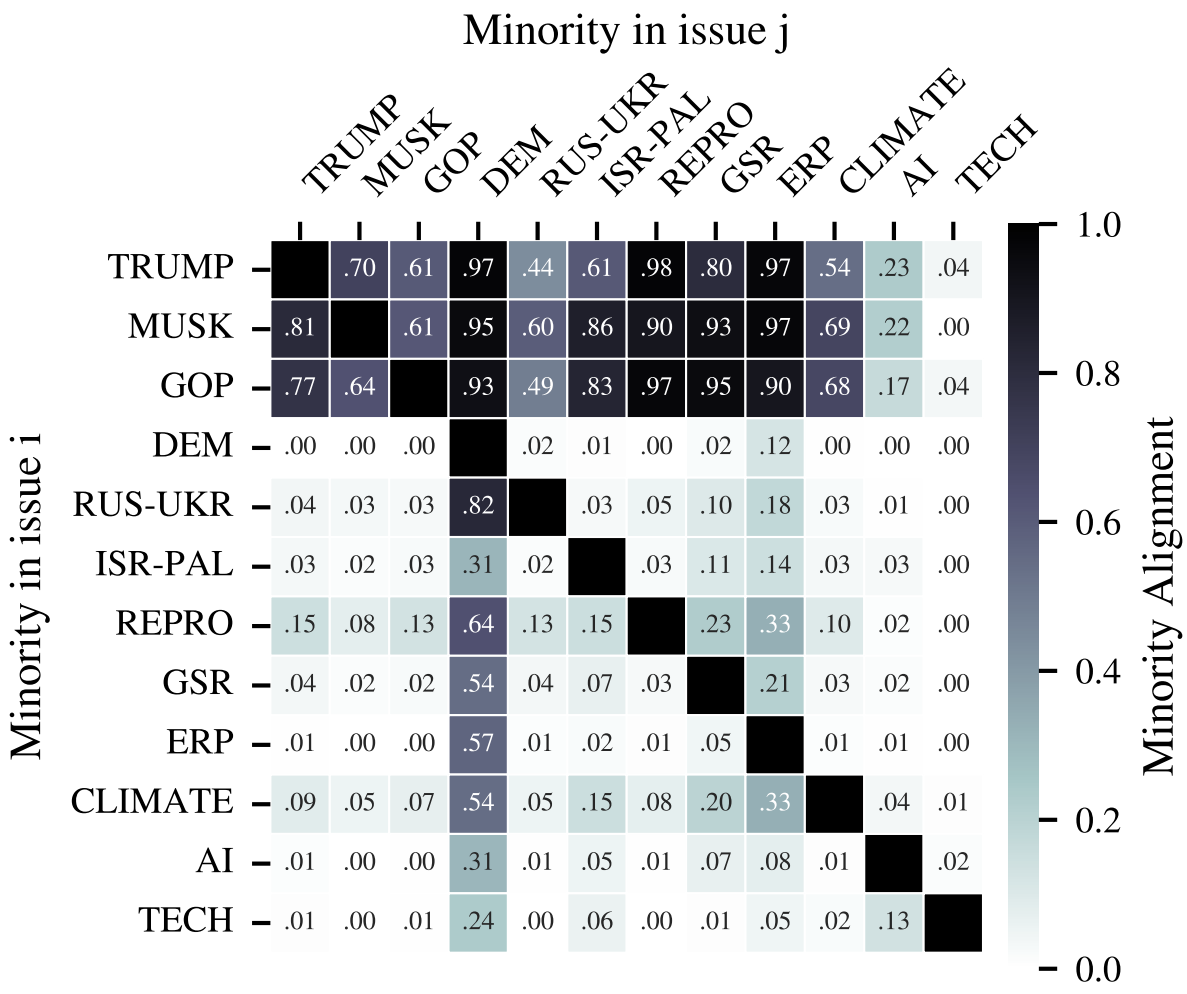


Figure 13: Minority alignment. Heatmap showing the conditional probability of alignment between minority opinions across different topics. Each cell represents the conditional probability $P(\text{minority}_j \mid \text{minority}_i)$: the probability that an individual belongs to the minority group in issue j (column) given that they belong to the minority group in issue i (row).

Supplementary Files

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

- anatomyplatformssupplementary.pdf