applied sciences

MDPI

Article

Modeling Narrative Activation and Affective Feedback in Ideologically Structured Telegram Ecosystems

Artūras Serackis [1], Dalius Matuzevičius [1], Gabriela Vdoviak [1], Henrikas Giedra [1], Ervinas Gisleris [1], Tomyslav Sledevič [1,*], Evaldas Bružė [2], Raminta Matulytė [2], Giedrė Sabaliauskaitė [3], Richard Andrew Paskauskas [2] and Tomas Lavišius [2]

[1] Department of Electronic Systems, Vilnius Gediminas Technical University, Saulėtekio Ave. 11, LT-10223 Vilnius, Lithuania; arturas.serackis@vilniustech.lt (A.S.); dalius.matuzevicius@vilniustech.lt (D.M.); gabriela.vdoviak@vilniustech.lt (G.V.); henrikas.giedra@vilniustech.lt (H.G.); ervinas.gisleris@vilniustech.lt (E.G.)

[2] Security Research Laboratory, Mykolas Romeris University, LT-08303 Vilnius, Lithuania

[3] Faculty of Public Governance and Business, Mykolas Romeris University, LT-08303 Vilnius, Lithuania

* Correspondence: tomyslav.sledevic@vilniustech.lt

Abstract

Digital social platforms have transformed political discourse into complex socio-technical environments characterized by rapid narrative diffusion, emotional amplification, and large-scale audience interaction. Understanding how sentiment and semantic alignment interact within such environments is important for analyzing polarization and patterns of audience response. This study examines narrative–audience interaction in Telegram political ecosystems using a combination of sentiment analysis, semantic similarity measures, and engagement metrics. Transformer-based language models are applied to quantify relationships between source posts and user-generated comments, enabling joint analysis of affective tone and topical alignment. The results reveal a consistent affective–semantic asymmetry: user responses tend to remain semantically aligned with source narratives while shifting toward more negative sentiment. This pattern indicates that disagreement is predominantly expressed through affective reframing rather than through divergence from the original topic. Further analysis shows systematic differences across ideological groups. Pro-government channels exhibit higher reach and more stable discourse alignment, while pro-opposition channels generate stronger engagement and more pronounced negative sentiment shifts. Neutral channels display intermediate characteristics. These findings demonstrate that online political discourse in Telegram is characterized by stable topical anchoring combined with systematic variation in emotional response.

Keywords: sentiment shift; semantic similarity; narrative classification; transformer-based models

Academic Editor: Rui Araújo

Received: 4 March 2026

Revised: 20 April 2026

Accepted: 21 April 2026

Published: 23 April 2026

1. Introduction

The rapid evolution of digital communication platforms has transformed political discourse into a complex, data-rich, and dynamically interacting ecosystem. Social media platforms enable large-scale narrative diffusion, emotional contagion, and engagement-driven amplification of political content. Systematic reviews demonstrate that online polarization, intergroup hostility, and ideological segregation have intensified in digitally mediated environments [1,2]. At the same time, platform dynamics tend to amplify divisive or emotionally charged content [3,4], reshaping how narratives spread and stabilize.

Although narrative diffusion processes in social networks have been extensively studied [5], understanding how narrative content and audience response interact remains a central challenge. Given the societal implications of disinformation and coordinated framing strategies [6,7], there is a growing need for empirical analyses that capture the interplay between semantic structure, emotional tone, and audience engagement.

Despite substantial progress in sentiment analysis, semantic modeling, and opinion dynamics, these approaches are often applied in isolation. Sentiment analysis has matured significantly, encompassing lexicon-based, machine learning, and transformer-based techniques [8], while emotion detection in political discourse has been operationalized through automated transcript analysis [9] and visualization frameworks [10]. At the same time, research on opinion dynamics has developed through analytical, simulation-based, and control-theoretic approaches [11], often abstracting away from the semantic structure of narratives. Agent-based models capture interaction effects and polarization mechanisms [12,13], yet typically rely on simplified representations of textual content. As a result, there remains a lack of empirical studies that jointly examine semantic alignment, sentiment dynamics, and engagement patterns within real-world communication environments.

The importance of such integrated analysis is particularly evident in platform-specific contexts. Telegram has emerged as a distinctive environment for political communication due to its broadcast-oriented structure, large channels, and relatively limited moderation [14]. Prior work documents the presence of collective narratives and ideological consolidation in Telegram communities [15], while studies of social learning and cascading activation highlight the active role of audiences in shaping discourse [16,17]. Cross-platform analyses further reveal consistent patterns of toxicity and engagement in online political communication [4], and research on digital media consumption shows how narrative framing interacts with audience response [18]. These findings motivate a closer empirical examination of how narrative structure and audience feedback co-evolve in Telegram ecosystems.

Existing methods provide important analytical tools but remain fragmented. Network-based approaches identify narrative structures and semantic centrality [19,20], while deep learning models support predictive analysis using sentiment and textual features [21,22]. Dynamical perspectives have been applied in related domains [23,24], and psychological research emphasizes the role of emotional coordination in communication [25,26]. However, empirical applications that combine semantic similarity, sentiment analysis, and engagement metrics in a unified analysis of real-world discourse remain limited. This fragmentation makes it difficult to capture the joint dynamics of narrative framing and audience response.

To address this gap, this study presents a large-scale empirical analysis of Telegram political discourse by integrating sentiment analysis, semantic similarity modeling, and engagement metrics into a unified analytical pipeline. Rather than introducing a new standalone methodological framework, the study combines established computational techniques to enable a systematic examination of narrative–audience interaction. This integrated approach allows for the joint analysis of affective tone, topical alignment, and participation intensity, providing a multidimensional perspective on online discourse dynamics.

The main contributions of this work are as follows:

- A comprehensive empirical analysis of narrative dominance, engagement intensity, and sentiment dynamics across ideologically distinct Telegram ecosystems.
- Identification of a systematic affective–semantic asymmetry, where audience responses remain topically aligned with source content while shifting toward more negative sentiment.

- Characterization of engagement patterns, showing that negative affect is a stronger driver of audience interaction than positive framing across all narrative groups.
- Comparative analysis of pro-government, pro-opposition, and neutral channels, revealing structurally different patterns of audience feedback and discourse alignment.

By providing an integrated empirical perspective on narrative structure, sentiment dynamics, and audience engagement, this study contributes to a better understanding of how political discourse evolves in digitally mediated environments and offers a scalable analytical approach for examining ideological communication ecosystems.

2. Related Work

The study of online discourse has been widely explored across multiple research domains, including sentiment and semantic analysis, narrative and polarization modeling, and engagement dynamics; while these strands provide valuable insights into specific aspects of digital communication, they are often examined separately, limiting a unified understanding of how content and audience interact in real-world environments.

2.1. Semantic and Sentiment Analysis in Online Communication

Recent advances in natural language processing have enabled increasingly accurate modeling of textual meaning and affect. Embedding-based methods allow quantification of semantic similarity and topical alignment between texts, supporting analysis of discourse coherence and thematic structure [24]. In parallel, automated sentiment analysis has been widely applied to large-scale online data to capture affective patterns in user-generated content [9]; while these approaches provide strong analytical foundations, they are typically applied independently, focusing either on semantic structure or emotional tone rather than their joint interaction.

2.2. Narrative Formation and Polarization

A substantial body of research examines how narratives emerge, evolve, and shape collective behavior in online environments. Prior work has explored narrative structures and their role in framing public discourse [5,19], as well as processes of group formation, ideological alignment, and polarization [1]. Computational approaches have enabled large-scale analysis of narrative diffusion and framing patterns [14]. However, much of this work emphasizes identifying narrative structures, with comparatively less attention to how these narratives are received and reinterpreted by audiences.

Recent research has highlighted the central role of political narratives as mechanisms for framing social reality, shaping collective interpretation, and reinforcing ideological alignment. Narratives function as interpretative structures through which actors construct meaning, assign responsibility, and position themselves within political discourse [19]. In digitally mediated environments, these narratives are often intertwined with emotionally charged and antagonistic language, which contributes to the intensification of polarization, the erosion of institutional trust, and the emergence of self-reinforcing information environments [1,27]. Empirical studies further show that political communication in online platforms is characterized by affective asymmetries, including the prevalence of negative or toxic interactions, particularly across ideological divides [4], and by engagement-driven dynamics that amplify emotionally salient or conflict-oriented content [2].

In the context of Telegram and similar platforms, prior work emphasizes the importance of structural and affordance-related factors in shaping narrative dissemination. Telegram channels enable large-scale broadcasting while supporting loosely connected communities, facilitating the spread of disinformation and coordinated messaging across networked clusters [28]. Studies also document the prominence of specific narrative forms,

https://doi.org/10.3390/app16094154

such as collective victimhood framing in extremist communities, which can intensify group identity and contribute to radicalization processes [15]. Moreover, platform-specific affordances give rise to diverse narrative profiles, ranging from coherent storytelling to fragmented or loosely connected narrative streams, depending on the thematic orientation and community structure [14]. These findings highlight that political narratives in online ecosystems are shaped not only by content, but also by platform architecture and interaction patterns.

2.3. Engagement Dynamics and Information Propagation

Another line of research focuses on how users interact with content and how information spreads across platforms. Engagement-driven mechanisms have been shown to influence content visibility and amplification, particularly in social media environments [3]. Studies of collective behavior further demonstrate how user interactions contribute to large-scale communication patterns [15]. Additional work highlights the relationship between engagement, toxicity, and polarization in online discourse [4,8]; while these studies capture important aspects of interaction, engagement is often treated as an outcome rather than as part of a coupled content–audience process.

2.4. Integrated Perspectives on Online Discourse

More recent research has begun to explore integrated perspectives that combine multiple dimensions of online communication. Agent-based and computational models simulate interactions between users and narratives to study emergent behavior in complex systems [13]. These approaches highlight the importance of considering feedback mechanisms, but often rely on stylized assumptions or focus on specific components of the system.

Building on these strands of research, this study combines sentiment analysis, semantic similarity, and engagement metrics to examine how narrative content and audience response interact in an empirical setting. Rather than introducing a new modeling framework, the analysis applies established computational methods in a unified manner to investigate their joint behavior in Telegram political discourse. This perspective enables a more comprehensive interpretation of how affective tone, topical alignment, and engagement co-evolve in practice. Although the analysis is conducted on Telegram data, the combination of methods used in this study can be applied to other online communication environments, supporting comparative investigations of narrative dynamics across platforms.

3. Analytical Approach

3.1. Conceptual Overview

This study adopts an integrated analytical approach to examine the interaction between narrative content and audience response in online discussions. The aim is to capture how the affective tone and semantic properties of textual content relate to audience engagement, and how these elements jointly characterize discourse dynamics in real-world communication environments.

At a high level, the analysis considers three interconnected components: (i) the narrative signal, representing the sentiment expressed in textual content; (ii) semantic alignment, capturing the topical similarity between posts and associated responses; and (iii) the engagement response, reflecting the intensity of audience interaction. These components are analyzed jointly to provide a multidimensional view of how content and audience behavior relate within the same communication instance.

To support this analysis, the interaction between posts and responses is interpreted from a dynamical perspective, where observed variables reflect different aspects of an evolving discourse process. In this setting, sentiment captures affective orientation, semantic similarity represents topical coherence, and engagement reflects the strength of audience response. This perspective enables the examination of narrative–audience interaction beyond static, single-variable analysis.

The role of feedback is particularly important in this context. In online environments, highly engaging content can amplify specific narratives, reinforce emotional framing, or contribute to shifts in discourse patterns. By jointly analyzing content characteristics and audience responses, the approach captures the recursive nature of communication, where audience reactions are closely linked to the properties of the originating content.

Although the analysis is applied in this study to Telegram political discourse, the same set of analytical components can be used to examine other online communication environments. The operational definitions of sentiment, semantic similarity, and engagement may vary depending on the dataset, but the overall structure of the analysis remains consistent, allowing for flexible application across different domains.

3.2. System Architecture and Dynamical Modeling

The interaction between content producers and their audiences is analyzed using a dynamical perspective that captures the relationship between narrative characteristics and audience response. The structure illustrated in Figure 1 represents this interaction as a feedback-oriented process, where audience engagement is treated as a response signal associated with previously published content. Rather than assuming a purely linear dissemination process, this representation highlights the interdependence between content properties and audience behavior.

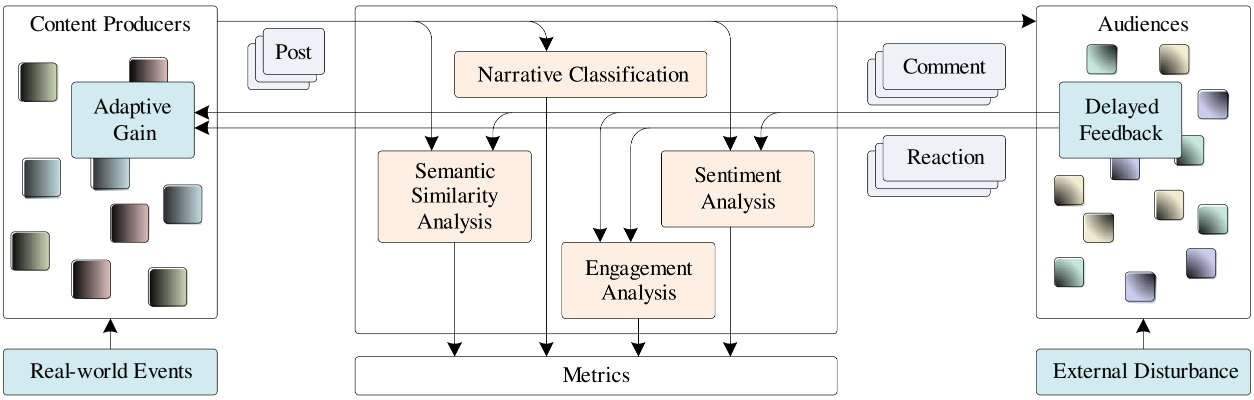

Figure 1. Structural representation of narrative–audience interaction, illustrating the relationship between content characteristics and audience response.

At a conceptual level, the input corresponds to the generation of narrative content, influenced by prior discourse and external events. The observed signal is characterized by both narrative intensity and sentiment, which together describe the affective and expressive properties of the content. Audience responses, in turn, provide measurable indicators of how content is received, including both engagement levels and sentiment expressed in reactions.

To support the analysis, the narrative–audience interaction is represented using a discrete-time state description. The state at time t is defined as:

$$x_t = \left[S_{\text{post},t}, S_{\text{response},t}, \bar{S}_{\text{responses},t}, S_{\text{sem},t}, K_{e,t}\right], \quad (1)$$

where $S_{\text{post},t}$ denotes the sentiment of the source content, $S_{\text{response},t}$ represents individual response sentiment, $\bar{S}_{\text{responses},t}$ is the aggregated audience sentiment, $S_{\text{sem},t}$ captures semantic alignment between content and responses, and $K_{e,t}$ is the normalized engagement intensity coefficient.

Within this representation, the forward relation captures how content characteristics are associated with audience response. After publication, the narrative becomes visible within the platform, where its reach is shaped by platform mechanisms and user activity. The resulting audience response reflects both the properties of the content and contextual factors influencing user behavior.

Two aspects are particularly relevant for interpreting the observed dynamics. First, audience reactions are subject to temporal delay, as engagement typically unfolds over time following content publication. Second, external events may influence audience behavior independently of the specific content, introducing variability in the observed response patterns. Considering these factors provides a more comprehensive view of how narrative characteristics and audience engagement are related within the empirical data.

3.3. Methodological Components

The analysis employs a multi-stage natural language processing pipeline to extract the variables used to characterize narrative content and audience response. The methodology focuses on quantifying narrative intensity, semantic alignment, and affective response, which together provide a structured representation of discourse dynamics.

Three transformer-based models are used to operationalize these components:

- Narrative classification was performed using the *MoritzLaurer/mDeBERTa-v3-base-mnli-xnli* model. This is a multilingual DeBERTa-v3 transformer fine-tuned on the MultiNLI and XNLI natural language inference (NLI) benchmarks spanning 100+ languages. It operates in a zero-shot fashion: given a text and a set of candidate labels phrased as hypotheses, it estimates the entailment probability for each hypothesis without task-specific fine-tuning. Its multilingual NLI pretraining makes it particularly suitable for cross-lingual political discourse where annotated corpora are scarce.
- Semantic similarity between posts and audience comments was computed with the *sentence-transformers/paraphrase-multilingual-mpnet-base-v2* model. This is a multilingual sentence-transformer based on the MPNet architecture, trained via contrastive learning on over one billion paraphrase pairs in 50+ languages. It maps input texts into a shared 768-dimensional dense vector space in which semantically related sentences are geometrically close. Cosine similarity between the post embedding and the mean comment embedding then quantifies discourse alignment–the degree to which audience discussion remains topically coherent with the original post. This model was selected for its strong cross-lingual transfer and robustness to informal social-media register.
- Sentiment analysis was conducted using the *cardiffnlp/twitter-xlm-roberta-base-sentiment* model. This is an XLM-RoBERTa transformer fine-tuned on the TweetEval benchmark for three-class sentiment classification (positive, neutral, negative). It is specifically optimized for short, informal social-media text and supports multiple languages. The model outputs class probabilities, from which a continuous sentiment score S_i is derived (Equation (3)). It outputs class probabilities that are subsequently transformed into a continuous sentiment score, enabling fine-grained analysis of affective dynamics.

The primary input to the analysis is the narrative intensity score derived from classification outputs. For each textual unit, the model produces a probability distribution over N candidate labels. To reduce ambiguity, a confidence threshold ($T = 0.5$) is applied. Only labels exceeding this threshold are considered active, while lower-confidence outputs are treated as unclassified:

$$D_i = \begin{cases} \text{label}_i & \text{if } P(\text{label}_i) \geq 0.5, \\ \text{Unclassified} & \text{if } P(\text{label}_i) < 0.5. \end{cases} \quad (2)$$

This thresholding procedure provides a consistent way to identify dominant narratives while filtering weak or ambiguous signals.

To ensure consistent polarity measurement, sentiment is represented as a continuous scalar derived from model-estimated class probabilities. For each textual unit i, the sentiment score is computed as:

$$S_i = P_i^{(+)} - P_i^{(-)}, \quad (3)$$

where $P_i^{(+)}$ and $P_i^{(-)}$ denote the predicted probabilities of positive and negative sentiment, respectively. The resulting score is bounded within $S_i \in [-1, 1]$, where values near $+1$ indicate strong positive polarity, values near -1 indicate strong negative polarity, and values near 0 correspond to neutral sentiment.

For aggregated analysis, sentiment values of associated responses are averaged:

$$\bar{S}_{\text{responses}} = \frac{1}{N} \sum_{j=1}^{N} S_j, \quad (4)$$

where N denotes the number of responses linked to a given content item. This provides a normalized representation of collective audience sentiment.

To examine the relationship between content and audience response, several derived measures are computed. Sentiment shift (ΔS) captures the difference between the sentiment of the source content and the aggregated audience response:

$$\Delta S = \bar{S}_{\text{responses}} - S_{\text{post}}, \quad (5)$$

where $\bar{S}_{\text{responses}}$ is the mean sentiment of the responses and S_{post} is the sentiment of the source content. This measure reflects how audience reactions differ in affective tone from the originating message.

Semantic alignment is computed as the cosine similarity between the embedding vector of the source content, v_p, and the mean embedding vector of associated responses, v_c:

$$S_{\text{sem}} = \frac{v_p \cdot v_c}{\|v_p\| \, \|v_c\|}. \quad (6)$$

Values of S_{sem} approaching 1 indicate strong topical coherence, while lower values reflect increasing divergence from the original narrative context.

To quantify the strength of audience response relative to exposure, an engagement intensity coefficient is defined as:

$$K_e = \frac{C}{V} \times 1000, \quad (7)$$

where C denotes the number of responses and V represents the exposure level of the content. This normalization enables comparison across content items with different reach. Higher values of K_e indicate stronger engagement relative to exposure.

The resulting representation defines a joint feature space in which narrative structure, affective response, semantic alignment, and engagement can be analyzed together.

https://doi.org/10.3390/app16094154

This allows for a systematic empirical examination of how different types of content are associated with distinct patterns of audience behavior.

4. Materials and Methods

This section describes the empirical dataset and the procedures used to implement the analytical approach in the context of the studied communication environment. We first outline the criteria for channel selection and dataset construction, ensuring a structured representation of ideologically distinct narrative ecosystems. We then detail the extraction of sentiment, semantic alignment, narrative intensity, and engagement variables from the data. Together, these elements provide the empirical basis for the representation introduced in Section 3, enabling quantitative analysis of narrative characteristics and audience response patterns.

4.1. Dataset and Channel Selection

The primary dataset comprises a curated sample of 88 high-visibility Telegram channels selected to capture the ideological structure of the Russian-language political information environment. Channel inclusion was based on sustained posting activity, audience reach, and relevance to public political discourse. The collected data span the period from 21 November 2025 to 20 January 2026, enabling temporal analysis of narrative activity and audience engagement during a continuous two-month observation window.

To situate the dataset, a brief contextual note is necessary. Following the onset of the full-scale Russian invasion of Ukraine in February 2022, Telegram emerged as one of the primary arenas for Russian-language political discourse: traditional media faced increasing censorship, while Telegram offered pseudonymous publishing, large broadcast channels, and minimal moderation. This created a highly polarized information ecosystem in which three broadly distinct orientations co-exist: pro-government channels aligned with official narratives and state-affiliated media frames; pro-opposition channels associated with independent journalists, exiled activists, or anti-war civic voices; and neutral channels that aggregate news or commentary without a consistently identifiable political alignment.

Each channel was assigned to one of three narrative groups according to its dominant editorial orientation: pro-government (Pro-Gov), pro-opposition (Pro-Opp), or neutral/informational. This grouping enables structured comparison of narrative framing, engagement patterns, and audience response across distinct political positions. As summarized in Table 1, the dataset covers channels with heterogeneous audience sizes and interaction intensities, including state-aligned sources, independent publishers, and opposition-oriented communities. This composition supports comparative analysis of how narrative positioning relates to engagement behavior and discourse alignment.

Table 1. Dataset statistics by narrative group.

Narrative Group	Channels	Posts	Comments Per Active Post	Availability (%)		Rate to Views (%)		Sentiment		Semantic Similarity
				Comment	Reaction	Comment	Reaction	Post	Comment	
Neutral	29	23,998	210	17.2	28.9	0.09	1.92	−0.08	−0.27	0.33
Pro-Gov	39	26,718	183	36.0	89.5	0.07	1.89	−0.08	−0.20	0.36
Pro-Opp	20	7995	355	15.7	56.5	0.14	2.20	−0.11	−0.22	0.36
Total (mean)	88	58,711	205	25.6	60.2	0.08	1.94	−0.09	−0.22	0.35

To capture interaction dynamics, the dataset integrates post-level content with engagement-related metadata, including reaction availability, comment status, and sentiment characteristics of both posts and responses. A substantial portion of the analyzed channels allows direct audience feedback; however, the availability of comments varies across the ideological spectrum, providing an additional dimension for distinguishing active public discourse from passive content consumption. The dataset includes aggregate indicators such as average reaction rate and average comment rate, together with sentiment estimates for both posts and associated comments. This multi-layered structure provides the variables required for the empirical analysis, enabling joint examination of narrative content, audience engagement, and discourse alignment.

To ensure robustness of semantic and sentiment inference, textual inputs were filtered prior to model processing. Entries were excluded if the text was empty, contained only whitespace, or consisted of fewer than three tokens. This filtering rule was applied uniformly across narrative classification, semantic similarity estimation, and sentiment analysis to reduce noise from non-informative or structurally incomplete content.

4.2. Implementation Details

This subsection describes how the analytical approach is applied to the empirical dataset; while Section 3 outlines the conceptual structure and computational components, the present subsection specifies their implementation in the context of Telegram-based discourse data. The dataset consists of posts and associated user comments collected from selected Telegram channels. Each post represents a unit of analysis, while the corresponding comment threads capture observable audience response. For each post–comment set, the variables defined in Section 3 are extracted and computed.

Textual data were preprocessed to ensure consistency across inputs. This included removal of non-textual artifacts, normalization of encoding, and segmentation of posts and comments into analyzable units. No language-specific filtering was imposed, allowing the multilingual models described in Section 3.3 to operate directly on the raw textual inputs.

For each post, narrative classification was performed to obtain a probability distribution over a predefined set of narrative labels. In this study, $N = 6$ narrative categories are considered: *Military Success*, *Civilian Welfare*, *Economic Hardship*, *Foreign Intervention*, *Internal Unrest*, and *Diplomatic Conflict*. These categories define the thematic space used in the empirical analysis. The dominant narrative was determined using a confidence threshold ($T = 0.5$), ensuring that only sufficiently strong signals are treated as active narrative assignments. Posts below this threshold were retained but treated as weakly classified instances.

Sentiment scores were computed independently for each post and for each associated comment. Comment-level sentiment values were then aggregated to obtain the mean audience sentiment, $\bar{S}_{\text{comments}}$, providing a collective representation of audience response. This aggregation enables direct comparison between the sentiment of the source content and the sentiment expressed in the discussion.

Semantic alignment between posts and their corresponding comment sets was computed using embedding-based cosine similarity. For each post, an embedding vector was generated and compared to the mean embedding of all associated comments. This provides a quantitative measure of discourse coherence within each post–comment interaction.

Engagement intensity was computed using the normalized comment rate per thousand views. The number of comments and views was extracted directly from the Telegram interface, enabling calculation of the engagement coefficient K_e. This normalization allows comparison across posts with different levels of exposure.

Temporal aspects of interaction are reflected implicitly in the data structure. Each post and its associated comments form a discrete observation in which audience response accumulates over time following publication; while explicit time-series modeling is not imposed, the aggregated variables capture the outcome of these temporal processes.

https://doi.org/10.3390/app16094154

The resulting dataset enables systematic empirical analysis of how narrative orientation relates to engagement patterns, audience sentiment, and discourse alignment. Each post is characterized by narrative probabilities, sentiment polarity, audience response metrics, and semantic similarity between posts and comments. Based on this representation, the study addresses four research questions: which narratives dominate the observed information space, which narratives are associated with stronger audience engagement, how narrative framing relates to audience sentiment response, and which narratives are associated with more aligned or more divergent discourse patterns.

5. Results and Discussion

The following section presents empirical findings addressing each research question through quantitative analysis of narrative prevalence, engagement behavior, sentiment response, and discourse alignment. All experiments were carried out on a workstation equipped with an NVIDIA GeForce RTX 4080 Super GPU with 16 GB VRAM. The software environment included Python 3.10.19, PyTorch 2.5.1, and CUDA 12.1.

5.1. Narrative Dominance and Engagement Intensity

Figure 2 presents the distribution of dominant narratives, defined as posts with a narrative probability score exceeding the 0.5 threshold, across the three political groups. The results reveal clear thematic asymmetries between groups. The pro-government group (blue) is strongly dominated by the *Foreign Intervention* narrative, which represents the most frequent category in the dataset. This indicates a pronounced emphasis on external actors and geopolitical framing as the central organizing theme. Additional prominence of *Military Success* and *Diplomatic Conflict* further supports an externally oriented narrative structure. In contrast, *Civilian Welfare* and *Internal Unrest* appear relatively infrequently within this group.

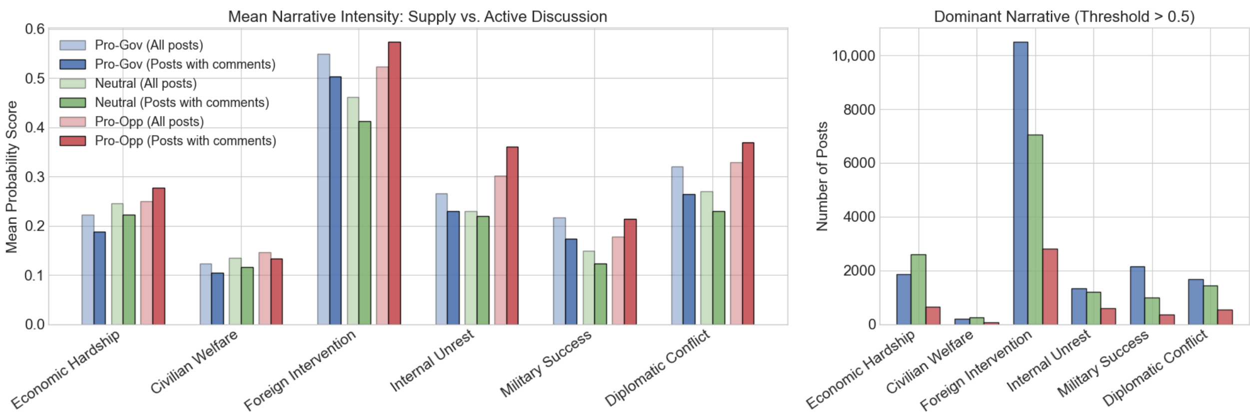

Figure 2. Narratives dominance in the information space.

The pro-opposition group (red) exhibits a more internally focused narrative profile; while *Foreign Intervention* remains present, its relative frequency is substantially lower than in the pro-government group. Instead, *Internal Unrest* and *Economic Hardship* occupy a more prominent role, indicating stronger attention to domestic instability and socio-economic conditions. The neutral group (green) combines elements of both patterns, with *Foreign Intervention* remaining highly frequent while other categories, such as *Economic Hardship*, *Diplomatic Conflict*, and *Military Success*, are more evenly distributed. Across all groups, *Civilian Welfare* remains consistently underrepresented, suggesting limited narrative centrality during the observed period.

Figure 3 compares engagement rate (comments per 1000 views), average comment count, and median views across narratives and ideological groups. The results show a clear separation between reach and interaction intensity. Pro-opposition channels consistently exhibit the highest engagement rates, particularly for *Internal Unrest*, where interaction levels substantially exceed those observed in other groups. Elevated engagement is also observed for *Military Success* and *Diplomatic Conflict*, indicating that conflict-related narratives are particularly effective in stimulating discussion within opposition-oriented communities.

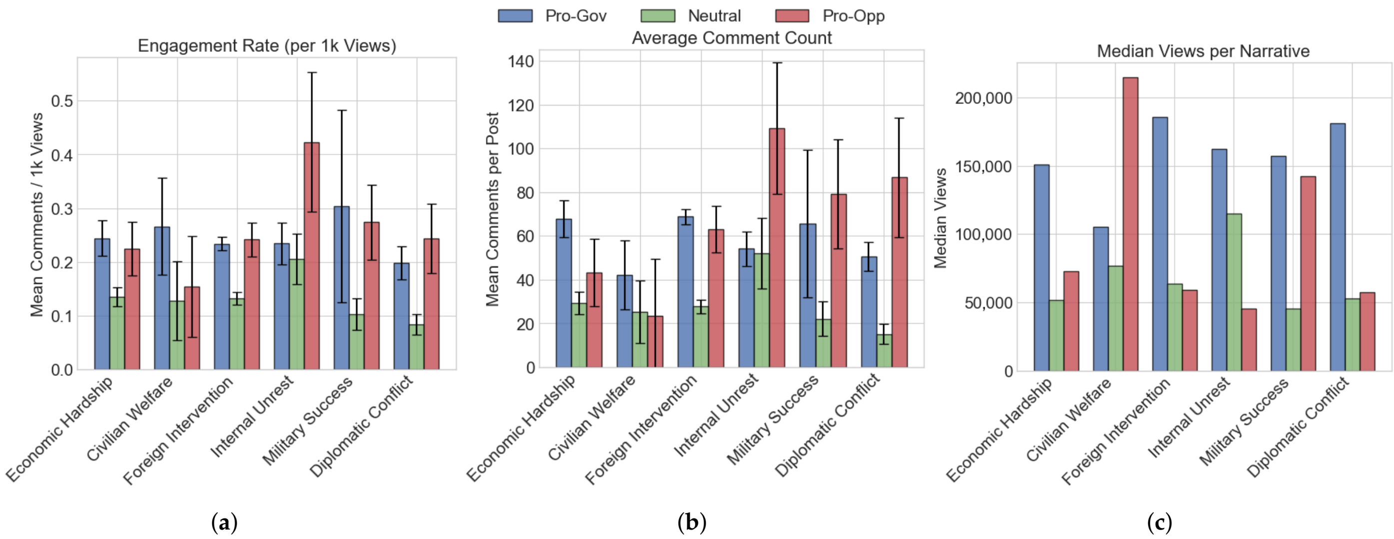

Figure 3. (**a**) Engagement rate by narrative, (**b**) comment activity by narrative, (**c**) post views by narrative.

In contrast, pro-government channels display more uniform engagement patterns across narrative categories, with moderate increases for *Military Success* and *Civilian Welfare*. This suggests a more stable interaction structure, with less pronounced engagement spikes. Neutral channels consistently show the lowest engagement rates across all categories, reflecting a predominantly informational communication style with limited audience activation.

A different pattern emerges when considering scale. Pro-government channels achieve the highest median views across most narrative categories, particularly in *Foreign Intervention*, *Diplomatic Conflict*, and *Military Success*, indicating strong broadcast reach. However, this visibility does not translate proportionally into interaction intensity. Pro-opposition channels, despite lower median views, generate comparatively high comment volumes, especially in *Internal Unrest* and *Diplomatic Conflict*. This divergence highlights a structural distinction between communication modes: pro-government discourse is characterized by high reach and lower interaction intensity, whereas pro-opposition discourse is more engagement-driven and participatory.

Figure 4 further illustrates the relationship between sentiment and audience response by showing the distribution of post sentiment (red) and associated comment volume (blue). Across all groups, posts are concentrated around neutral to mildly negative sentiment values, with peaks near zero. In contrast, comment activity is consistently skewed toward more negative sentiment ranges.

In the pro-government group (Figure 4a), comments exhibit a clear peak at moderately negative sentiment levels, despite posts being concentrated near neutrality. A similar pattern is observed in the neutral group (Figure 4b), where comment volume increases disproportionately in the negative range. The pro-opposition group (Figure 4c) displays an even stronger effect, with comment activity rising sharply as sentiment becomes more negative. Across all groups, strongly positive sentiment is associated with comparatively low comment volume.

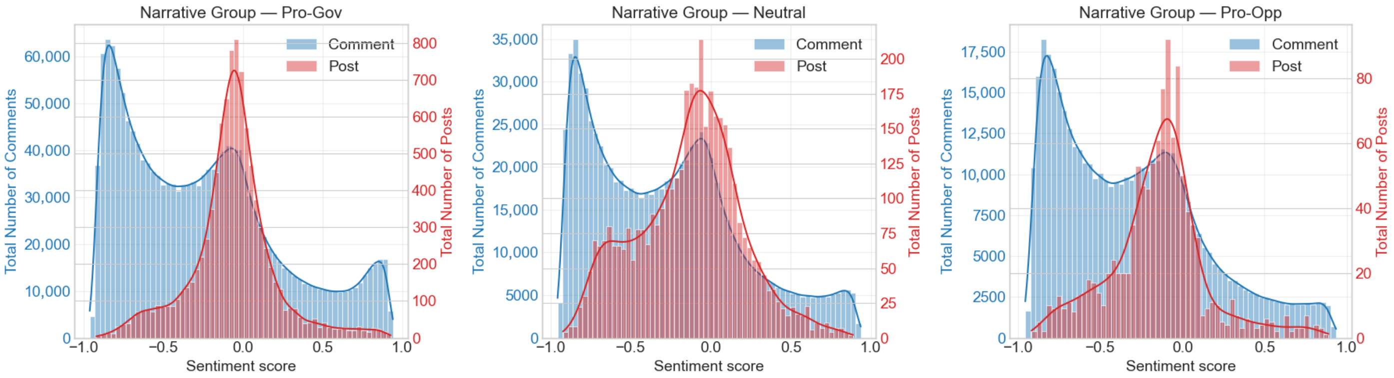

(**a**) (**b**) (**c**)

Figure 4. Sentiment vs. engagement for narrative groups: (**a**) Pro-Gov, (**b**) Neurtal, and (**c**) Pro-Opp.

These results indicate a consistent relationship between sentiment and engagement: negative affect is more strongly associated with audience interaction than neutral or positive framing. This pattern holds across ideological groups, suggesting that emotionally negative content plays a central role in driving participatory discourse within the observed Telegram communication environment.

5.2. Sentiment Feedback and Discourse Alignment

The joint distribution of comment sentiment and semantic similarity to the corresponding source post provides a detailed view of how emotional tone and topical alignment interact within audience responses. The density heatmap (Figure 5) shows that the highest concentration of comments occurs at moderate levels of semantic similarity (approximately 0.2–0.4), with a pronounced peak at slightly negative sentiment values (around −0.2 to 0). In addition, a distinct high-density region is observed at strongly negative sentiment levels (approximately −0.7 to −0.9). These patterns indicate that audience responses typically remain topically related to the original post while expressing predominantly negative evaluations.

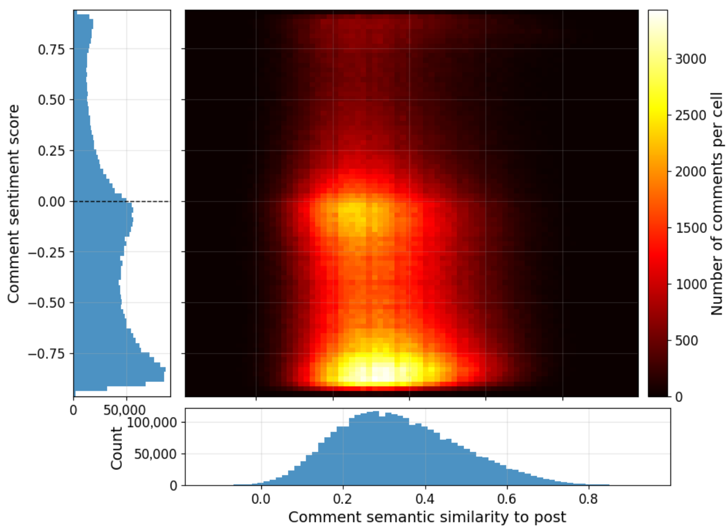

Figure 5. Comment sentiment score vs. comment semantic similarity to post.

The marginal distributions support this interpretation. Semantic similarity follows a unimodal distribution centered around mid-range values (approximately 0.3), suggesting that most comments maintain moderate alignment with the source content rather than diverging completely. In contrast, sentiment is clearly skewed toward negative polarity, with substantially higher density below zero. Strongly positive sentiment appears infrequently and does not form a significant concentration in the joint space. This asymmetry indicates that negative affect dominates audience responses even when topical alignment is preserved.

The group-level analysis reveals consistent but structurally distinct patterns across ideological categories. In the pro-government profile (Figure 6a), audience responses are systematically more negative than the originating posts across nearly all narrative categories. Posts associated with *Civilian Welfare* and *Military Success* exhibit positive sentiment, yet corresponding comments shift toward neutrality or negativity, resulting in consistently negative sentiment differences. For already negative narratives, such as *Economic Hardship*, *Internal Unrest*, and *Diplomatic Conflict*, comments further intensify negative sentiment. These patterns are reflected in the sentiment shift estimates, where mean differences remain below zero across categories. Correlation analysis shows that stronger emphasis on *Internal Unrest* and *Diplomatic Conflict* is associated with more negative audience sentiment, while *Military Success* exhibits a weak positive association.

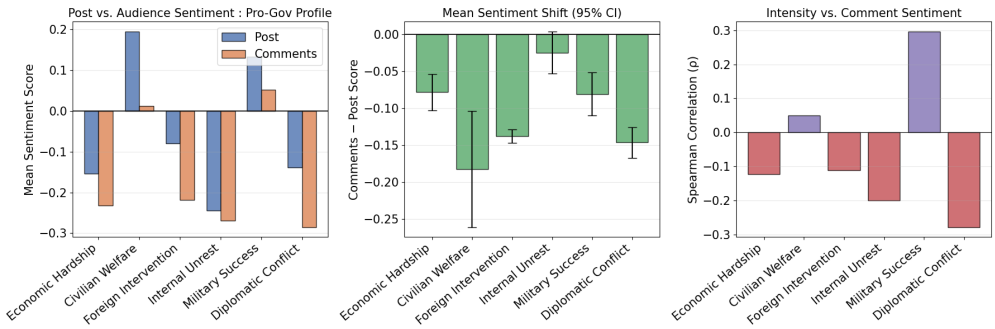

(**a**)

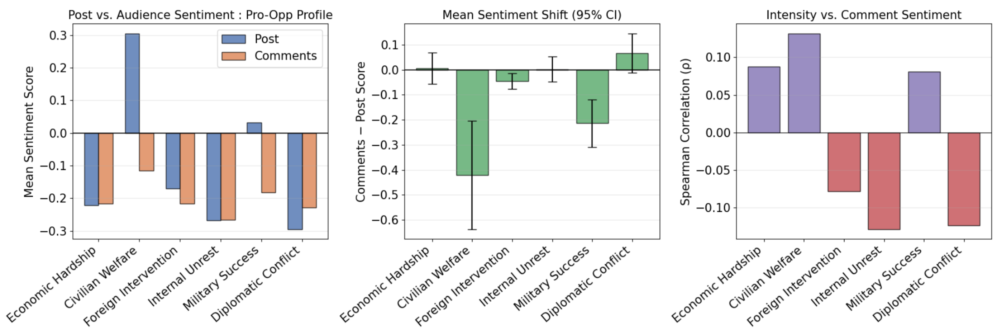

(**b**)

Figure 6. Post vs. audience sentiment, audience sentiment shift, narrative intensity vs. comment sentiment in (**a**) Pro-Gov and (**b**) Pro-Opp profile.

The pro-opposition profile (Figure 6b) displays a more heterogeneous response structure. For several narratives, particularly *Civilian Welfare* and *Military Success*, posts are relatively positive, but comments exhibit substantial downward shifts in sentiment. However, unlike the pro-government case, some narratives show weaker or even slightly positive shifts, indicating that sentiment response is more context-dependent. Correlation patterns further highlight this variation: narrative intensity is positively associated with comment sentiment for some categories (e.g., *Civilian Welfare*, *Economic Hardship*), but negatively associated for others (e.g., *Foreign Intervention*, *Internal Unrest*). This suggests a more polarized and narrative-contingent response structure.

Discourse alignment exhibits both stability and variation across groups. In the pro-government setting (Figure 7a), semantic similarity scores cluster tightly around the group mean (approximately 0.36), with substantial overlap across narrative categories. This indicates a relatively stable level of topical alignment between posts and comments. However, statistical tests identify selective differences, particularly for *Internal Unrest* and *Military Success*, suggesting that some narratives produce distinct alignment patterns. Correlation analysis shows that stronger emphasis on *Civilian Welfare* and *Military Success* is associated with slightly higher alignment, whereas *Foreign Intervention* exhibits a weak negative association.

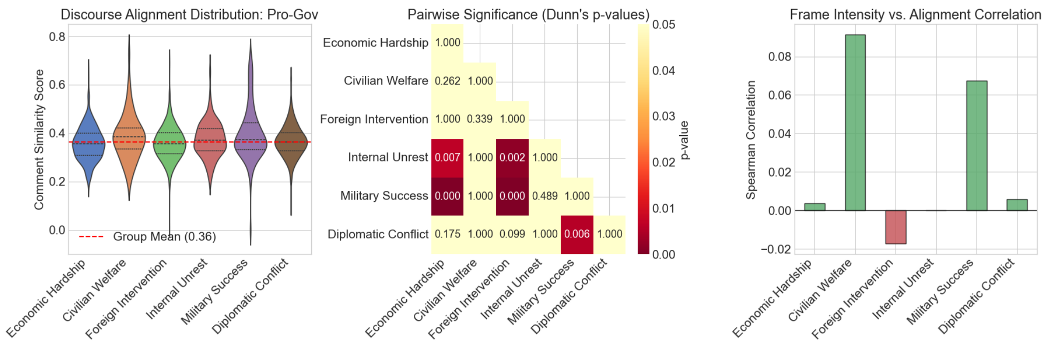

(**a**)

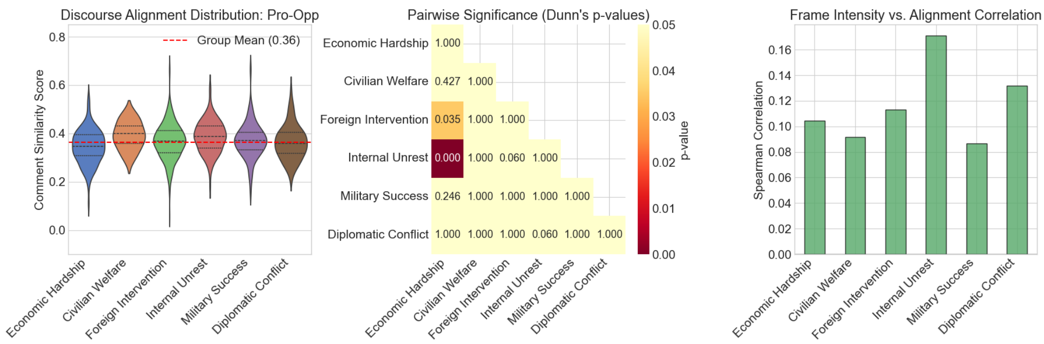

(**b**)

Figure 7. (**a**) Pro-Gov, (**b**) Pro-Opp narrative groups. Discourse alignment distribution, pairwise significance (Dunn's p-values), and intensity vs. alignment correlation.

The pro-opposition profile (Figure 7b) demonstrates greater variability in alignment patterns. Although the overall mean similarity is comparable to the pro-government group, the distribution is more dispersed across narratives. Statistical comparisons indicate more pronounced differences between categories, particularly involving *Economic Hardship* and *Internal Unrest*. In contrast to the mixed correlations observed in the pro-government group, all narrative-intensity correlations in the pro-opposition setting are positive, suggesting that increased narrative emphasis is consistently associated with stronger topical alignment. This indicates that discourse in pro-opposition channels is more sensitive to narrative framing.

Figure 8 highlights differences in the relationship between sentiment and engagement across groups. In the pro-government segment, post sentiment is moderately correlated with comment sentiment ($r \approx 0.35$), indicating partial alignment in emotional tone, while engagement metrics show weak relationships with sentiment. This suggests that emotional framing has limited influence on interaction intensity in this group.

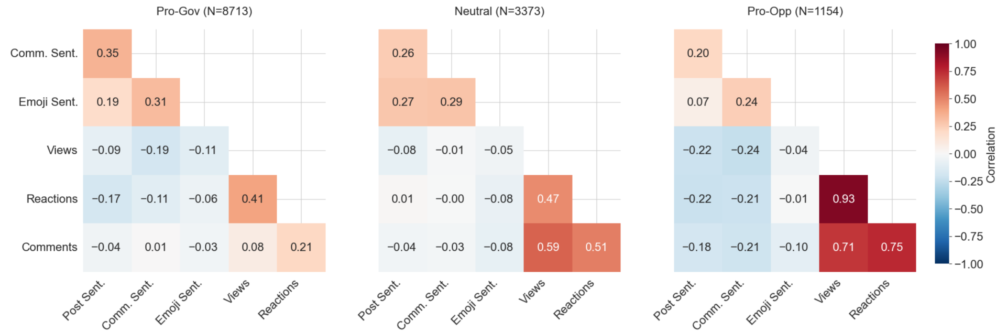

Figure 8. Sentiment and engagement correlation by narrative group.

The neutral group exhibits a different pattern, where engagement metrics are more strongly interrelated than sentiment variables. Views show moderate correlations with reactions and comments ($r \approx 0.47$ and $r \approx 0.59$), indicating that visibility is the primary driver of interaction. Sentiment remains weakly associated with engagement, suggesting that audience response is less influenced by emotional tone.

The pro-opposition group shows the strongest coupling between engagement metrics, with very high correlations between views, reactions, and comments (up to $r \approx 0.93$). At the same time, sentiment variables exhibit weak negative correlations with engagement, indicating that higher interaction levels are associated with more negative discourse. This pattern suggests that participation intensity is amplified independently of positive emotional alignment.

Figure 9 presents the temporal dynamics of semantic and emotional drift across narrative ecosystems over a 120-h window following publication. Figure 9a shows that semantic similarity does not collapse sharply but instead exhibits gradual divergence patterns across groups. The Pro-Gov narrative maintains the highest and most stable level of semantic alignment with the source post throughout the observation period, indicating sustained topic coherence. In contrast, the Pro-Opp and Neutral groups display greater volatility and episodic drops in similarity, particularly after the first 48–72 h, suggesting a higher propensity for discursive branching and contextual reframing. This pattern implies that topic drift is not purely a time-driven decay process but is structurally mediated by narrative community norms.

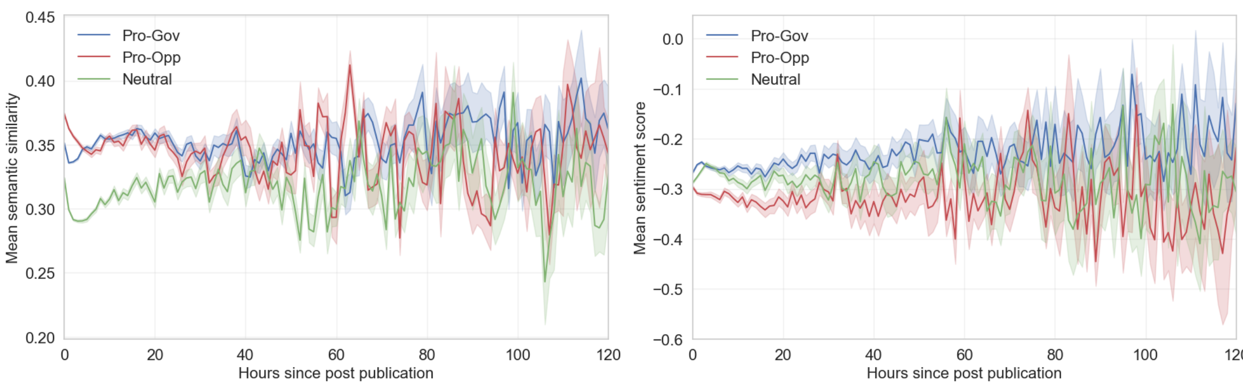

(**a**) (**b**)

Figure 9. Temporal dynamics of (**a**) semantic and (**b**) emotional drift across narrative ecosystems.

Figure 9b reveals even more evident divergence in emotional trajectories; while all groups begin with moderately negative mean sentiment, the Pro-Opp narrative exhibits increasing emotional volatility over time, with larger fluctuations and deeper negative excursions in later stages of the discussion cycle. The Pro-Gov ecosystem, by comparison, shows a more controlled and gradually stabilizing sentiment trend, whereas the Neutral group remains intermediate in both magnitude and variability. Importantly, emotional drift appears less monotonic than semantic drift, suggesting that affective amplification operates through episodic intensification rather than steady decay. The narrative ecosystems differ not only in topic persistence but also in the regulation and escalation of collective emotional expression over time.

5.3. Resistance Density Maps

The resistance density maps (Figure 10) provide a two-dimensional representation of audience response by jointly considering semantic alignment (x-axis) and sentiment shift relative to the source post (y-axis). Using the global mean semantic similarity and a zero sentiment-shift threshold as reference axes, the space is partitioned into four regions: active support (high alignment, positive shift), active resistance (high alignment, negative shift), topic hijacking (low alignment, positive shift), and topic divergence (low alignment, negative shift). This representation allows for a structured interpretation of how audience responses combine topical alignment with affective direction.

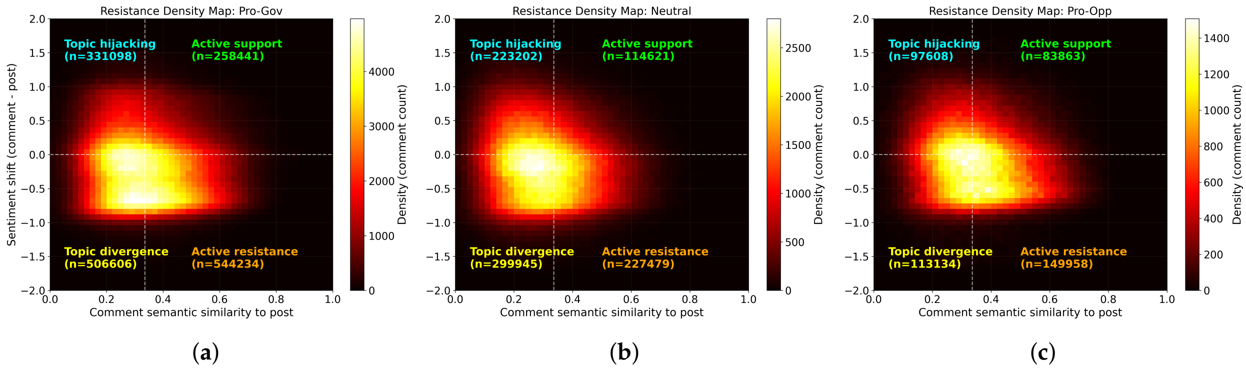

Figure 10. Resistance density maps illustrating the relationship between semantic similarity and sentiment shift across (**a**) pro-government, (**b**) neutral, and (**c**) pro-opposition narrative groups. Based on all observed comments.

Across all three ecosystems, the highest density of observations is concentrated near moderate semantic similarity and slightly negative sentiment shifts. This indicates that most comments remain loosely aligned with the source content while expressing mildly critical or corrective sentiment. Despite this common baseline, the distribution across regions differs systematically between groups.

In the pro-government ecosystem (Figure 10a), the largest concentrations are observed in the topic divergence (n = 506,606) and active resistance (n = 544,234) regions. This suggests that audience responses frequently combine moderate alignment with negative sentiment shifts. At the same time, substantial activity is also present in the active support (n = 258,441) and topic hijacking (n = 331,098) regions, indicating a heterogeneous interaction structure with both supportive and critical responses. The neutral group (Figure 10b) exhibits a more balanced distribution across regions. However, topic divergence (n = 299,945) remains more prominent than active support (n = 114,621), indicating that even in less polarized environments, audience responses tend to shift away from the original sentiment baseline while maintaining partial topical connection. The pro-opposition ecosystem (Figure 10c) shows the strongest concentration in resistance-oriented regions, particularly active resistance (n = 149,958), alongside notable density in topic divergence (n = 113,134). This pattern reflects a more consistently adversarial response structure, where audience reactions remain topically aligned but shift toward more negative sentiment.

While semantic alignment remains relatively stable across groups, sentiment shifts introduce clear asymmetries. Pro-opposition discussions exhibit denser clustering in negative-shift regions, neutral channels show more moderate dispersion, and pro-government channels combine high activity with a more heterogeneous distribution. These patterns indicate that differences between ecosystems are driven less by topical alignment than by the direction and intensity of sentiment expressed in responses.

Figure 11 presents topic-specific density maps focusing on highly engaged comments (top 10 per post with at least 100 reactions). These filtered representations highlight the most visible audience responses and reveal more pronounced structural differences between ecosystems. Across most narrative categories, pro-opposition discussions show a stronger concentration in the active resistance region, indicating that highly engaged comments tend to remain on-topic while expressing negative sentiment. In contrast, pro-government channels display a more mixed distribution between active support and active resistance, with greater dispersion toward topic divergence in some categories. Neutral channels again occupy an intermediate position, with lower density peaks and more symmetric distributions around the sentiment baseline.

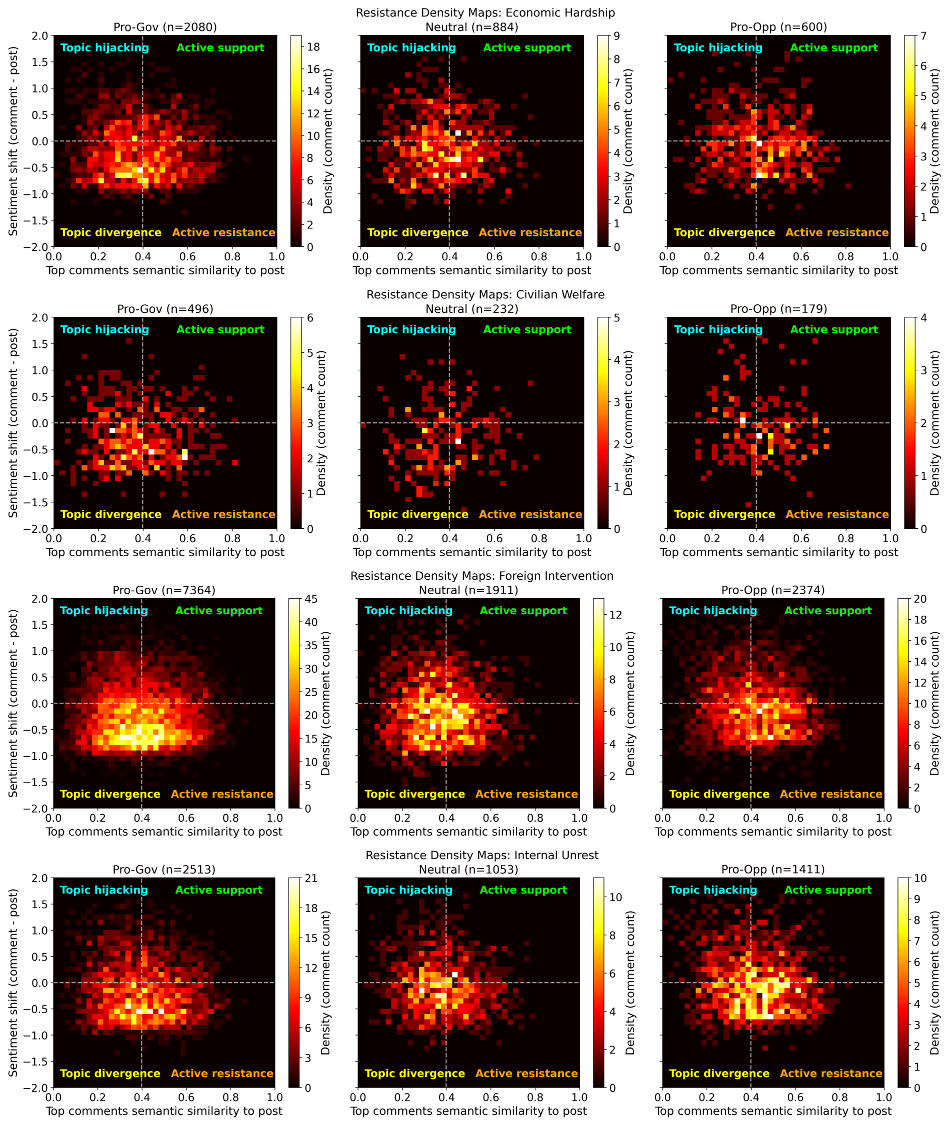

Figure 11. *Cont.*

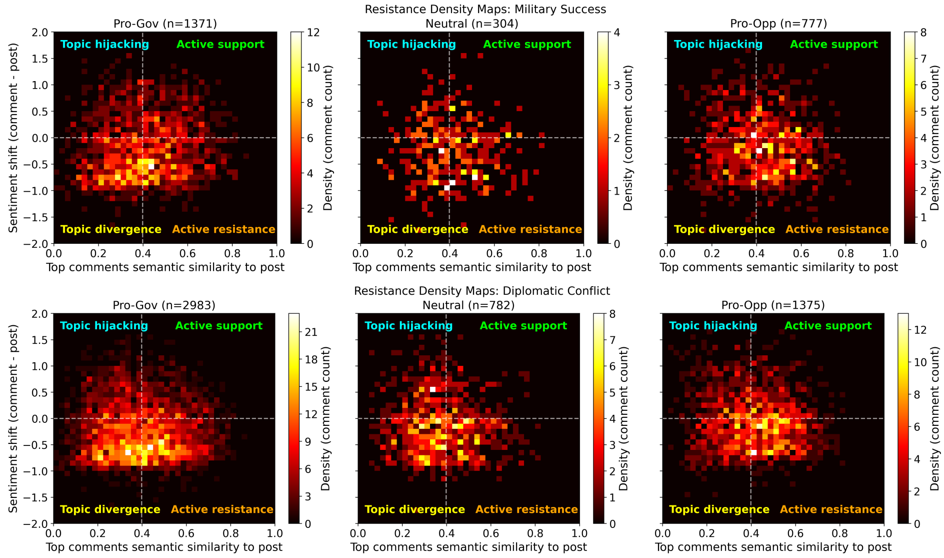

Figure 11. Topic-specific resistance density maps across narrative ecosystems (top engaged comments). Only the top 10 comments per post with a minimum of 100 reactions are included, and only posts with at least 100 total comments are considered. n denotes the number of comments represented in each subplot.

The differences are particularly pronounced for *Economic Hardship*, *Internal Unrest*, and *Diplomatic Conflict*. In these categories, pro-opposition discussions exhibit stronger clustering in the negative-shift region while maintaining moderate-to-high semantic similarity, indicating consistent critical engagement with the source content. Pro-government distributions, while also showing negative-shift mass, are more dispersed, suggesting a mixture of reinforcement and critique.

For *Foreign Intervention* and *Military Success*, both ecosystems maintain high semantic alignment, but differ in sentiment structure. Pro-opposition responses remain concentrated in negative-shift regions, whereas pro-government responses show broader variation, indicating coexistence of supportive and critical reactions. The *Civilian Welfare* category exhibits lower density and more diffuse patterns across all groups, suggesting weaker polarization and lower engagement intensity.

These results show that while topical alignment remains relatively stable across narrative ecosystems, sentiment direction varies systematically, particularly among highly engaged responses. This indicates that differences in audience behavior are driven primarily by affective framing rather than by divergence in topical focus.

5.4. Empirical Contributions and Insights

This study provides several empirical insights into the relationship between narrative framing and audience response in Telegram political discourse. By jointly analyzing sentiment, semantic alignment, and engagement patterns, the results reveal consistent structural properties of narrative–audience interaction across ideologically distinct communication environments.

First, the analysis identifies a systematic affective–semantic asymmetry in audience response; while comments remain strongly aligned with the semantic content of source posts, their sentiment consistently shifts toward more negative values. This pattern is observed across all ideological groups, indicating that audience discussion tends to preserve topical coherence while amplifying negative affect. This finding extends prior observations of emotional amplification in digital media [3,4] by demonstrating that negativity bias coexists with high semantic alignment in real-world discourse.

Second, the results show that negative sentiment is a strong and consistent driver of audience engagement. Posts associated with more negative affect generate higher levels of interaction relative to exposure, suggesting that emotional intensity plays a central role in shaping participation dynamics. This relationship holds across different channel types, indicating that negativity-driven engagement is a general feature of the observed communication environment rather than a group-specific effect.

Third, the study reveals clear structural differences between ideological groups in terms of audience interaction patterns. Pro-government channels are characterized by high visibility but comparatively lower engagement intensity, consistent with a broadcast-oriented communication structure. In contrast, pro-opposition channels exhibit lower reach but substantially higher engagement, reflecting more interactive and participatory discourse dynamics. These differences indicate that narrative positioning is associated with distinct modes of audience response, ranging from passive consumption to active discussion.

Finally, the joint analysis of sentiment and semantic alignment provides additional insight into discourse structure. The results show that high semantic coherence does not necessarily correspond to affective agreement; instead, audiences often engage critically with content while remaining within the same topical space. This decoupling of semantic alignment and sentiment response highlights the importance of considering multiple dimensions of discourse simultaneously, as single-variable analyses may overlook key aspects of audience behavior.

5.5. Limitations

Several limitations should be considered when interpreting the results. First, narrative identification is based on zero-shot classification using predefined frame descriptions; while this approach enables scalable multilingual analysis without manual annotation, the resulting labels reflect probabilistic semantic matching rather than ground-truth annotations. Accordingly, narrative assignments should be interpreted as approximate indicators of thematic emphasis rather than exact categorizations.

Second, sentiment and semantic alignment measures rely on multilingual transformer models. Although these models provide strong cross-lingual capabilities, their performance may vary across languages, stylistic registers, and platform-specific discourse conventions. In addition, the sentiment metric captures polarity rather than nuanced emotional states, while semantic similarity reflects textual coherence rather than agreement or endorsement. As a result, the analysis focuses on structural patterns in discourse rather than fully capturing subjective audience intent.

Third, the study is observational and limited to a single platform. Telegram-specific engagement indicators (views, reactions, comments) reflect the affordances of this environment and may not directly generalize to other social media platforms or broader populations. The findings therefore describe associations between narrative framing and audience response, without establishing causal relationships.

The set of narrative categories represents a bounded and simplified representation of political discourse. Alternative or more fine-grained labeling schemes could reveal additional structure and variation. Future work may extend this analysis through annotated datasets, cross-platform comparisons, and longitudinal designs to further examine the robustness of the observed patterns.

5.6. Legal and Regulatory Considerations for Automated Discourse Analysis

The observed platform-specific differences in discourse dynamics, particularly the presence of heavy negative sentiment tails and emotionally amplified engagement patterns, raise important considerations beyond purely computational analysis. From a regulatory perspective, systems designed to monitor and analyze online discussions operate within a complex legal environment shaped by European data protection and emerging artificial intelligence governance frameworks. In particular, the automated collection and processing of user-generated content, even when publicly available, may trigger obligations under the General Data Protection Regulation (GDPR), including requirements related to lawful basis, data minimization, and transparency. Furthermore, the increasing recognition of foreign information manipulation and interference (FIMI) as a security concern highlights the dual-use nature of such analytical systems, which may serve both research and policy-oriented monitoring purposes.

In this context, the development of computational tools for discourse analysis must consider principles such as privacy by design and by default, as well as forthcoming requirements under the European Union AI Act for high-risk AI systems. These include the need for human oversight, robustness against bias, and the implementation of data governance mechanisms that ensure accountability and traceability. Additional challenges arise in distinguishing between legitimate political expression and manipulative or disruptive behavior, as this boundary remains legally and ethically contested in European jurisprudence and policy guidance. Consequently, while the analytical approach applied in this study provides useful insights into discourse patterns, its application in operational settings would require careful alignment with legal standards, ethical safeguards, and ongoing regulatory developments.

6. Conclusions

The aim of this study was to examine narrative–audience interaction dynamics in Telegram political discourse by jointly analyzing sentiment, semantic alignment, and engagement patterns. The results suggest that, within the analyzed dataset, audience responses tend to remain semantically aligned with source posts while shifting toward more negative sentiment. This indicates that reactions are often expressed through affective modulation rather than through substantial divergence from the original narrative context.

This pattern is relevant for understanding how online political discourse evolves in practice. Rather than abandoning the original topic, users frequently engage within the same semantic frame while adjusting its emotional interpretation. In the analyzed Telegram channels, this results in a consistent tendency toward more critical or negative audience responses, even when the originating content is neutral or positive. These findings highlight the importance of considering both semantic coherence and affective dynamics when interpreting patterns of online engagement.

The analysis also reveals differences across the examined narrative ecosystems. Pro-government channels exhibit higher levels of content visibility and more stable semantic alignment, whereas pro-opposition channels show stronger engagement intensity and more pronounced negative sentiment shifts. Neutral channels display intermediate characteristics. These differences suggest that narrative orientation is associated not only with thematic emphasis but also with distinct patterns of audience interaction and feedback.

From a methodological perspective, the study demonstrates that combining sentiment analysis, semantic similarity estimation, and engagement metrics provides a coherent basis for examining discourse dynamics in an empirical setting. The joint analysis of these dimensions enables a more nuanced interpretation of how narrative content and audience response are related within the observed communication environment.

The findings should be interpreted in the context of the analyzed dataset and platform. The study focuses on a specific set of Telegram channels over a defined observation period, and the results may not generalize to other platforms or broader populations. Future work

https://doi.org/10.3390/app16094154

may extend this analysis to additional communication environments, incorporate longitudinal modeling, or explore alternative narrative representations to further investigate the dynamics of online discourse.

Author Contributions: Conceptualization, G.V. and T.S.; methodology, T.S., H.G. and A.S.; software, T.S., E.G. and G.V.; validation, A.S., R.M., T.L. and D.M.; formal analysis, E.G. and H.G.; investigation and analysis, R.M., D.M., E.B., G.S., R.A.P., T.L. and G.V.; data curation, H.G., G.S. and R.A.P.; writing—original draft preparation, T.S., A.S., G.V. and D.M.; writing—review and editing, A.S., E.G., H.G., E.B., R.A.P., R.M., T.L. and D.M.; visualization, G.V. and E.G. All authors have read and agreed to the published version of the manuscript.

Funding: Research was conducted as part of the execution of project titled "Mission-driven Implementation of Science and Innovation Programmes" (No. 02-002-P-0001), funded by the Economic Revitalization and Resilience Enhancement Plan "New Generation Lithuania".

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original code and data presented in this study are openly available in Zenodo at https://doi.org/10.5281/zenodo.18846833.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bliuc, A.M.; Bouguettaya, A.; Felise, K.D. Online intergroup polarization across political fault lines: An integrative review. *Front. Psychol.* **2021**, *12*, 641215. [CrossRef]
2. Arora, S.D.; Singh, G.P.; Chakraborty, A.; Maity, M. Polarization and social media: A systematic review and research agenda. *Technol. Forecast. Soc. Change* **2022**, *183*, 121942. [CrossRef]
3. Milli, S.; Carroll, M.; Wang, Y.; Pandey, S.; Zhao, S.; Dragan, A.D. Engagement, user satisfaction, and the amplification of divisive content on social media. *PNAS Nexus* **2025**, *4*, pgaf062. [CrossRef] [PubMed]
4. Falkenberg, M.; Zollo, F.; Quattrociocchi, W.; Pfeffer, J.; Baronchelli, A. Patterns of partisan toxicity and engagement reveal the common structure of online political communication across countries. *Nat. Commun.* **2024**, *15*, 9560. [CrossRef] [PubMed]
5. Gurung, M.I.; Agarwal, N. Narrative diffusion in social networks: A survey. *Soc. Netw. Anal. Min.* **2025**, *16*, 19. [CrossRef]
6. Recuero, R. A systemic framework for disinformation on social media platforms. *Platforms Soc.* **2025**, *2*, 29768624251367199. [CrossRef]
7. Entman, R.M.; Usher, N. Framing in a fractured democracy: Impacts of digital technology on ideology, power and cascading network activation. *J. Commun.* **2018**, *68*, 298–308. [CrossRef]
8. Anderson, T.; Sarkar, S.; Kelley, R. Analyzing public sentiment on sustainability: A comprehensive review and application of sentiment analysis techniques. *Nat. Lang. Process. J.* **2024**, *8*, 100097. [CrossRef]
9. Cochrane, C.; Rheault, L.; Godbout, J.F.; Whyte, T.; Wong, M.W.C.; Borwein, S. The automatic analysis of emotion in political speech based on transcripts. *Political Commun.* **2022**, *39*, 98–121. [CrossRef]
10. de Diego, I.M.; Fernández-Isabel, A.; Ortega, F.; Moguerza, J.M. A visual framework for dynamic emotional web analysis. *Knowl.-Based Syst.* **2018**, *145*, 264–273. [CrossRef]
11. Zareer, M.; Selmic, R.R. A Survey on Opinion Dynamics in Social Media Networks: Analysis, Simulation, and Control. *IEEE Trans. Comput. Soc. Syst.* **2025**, *13*, 2626–2660. [CrossRef]
12. Zhang, K.; Dong, C.; Guo, Y.; Zhou, W.; Yu, G.; Mi, J. Lagged Stance Interactions and Counter-Spiral of Silence: A Data-Driven Analysis and Agent-Based Modeling of Technical Public Opinion Events. *Systems* **2025**, *13*, 417. [CrossRef]
13. Bojić, L.; Ilić, V.; Prodanović, V.; Vuković, V. An Agent-Based Simulation of Politicized Topics Using Large Language Models: Algorithmic Personalization and Polarization on Social Media. *Chin. Political Sci. Rev.* **2025**, 1–31. [CrossRef]
14. Willaert, T. A computational analysis of Telegram's narrative affordances. *PLoS ONE* **2023**, *18*, e0293508. [CrossRef] [PubMed]
15. Boussalis, C.; Craig, C.; Rudkin, A.F. Collective victimhood narratives in far-right communities on Telegram. *Inf. Commun. Soc.* **2025**, *28*, 1407–1424. [CrossRef]
16. Brady, W.J.; Jackson, J.C.; Lindström, B.; Crockett, M. Algorithm-mediated social learning in online social networks. *Trends Cogn. Sci.* **2023**, *27*, 947–960. [CrossRef]
17. Landis, B.T.; Allen, W.L. Cascading activation revisited: How audiences contribute to news agendas using social media. *Digit. J.* **2022**, *10*, 537–555. [CrossRef]

18. Famulari, U.; Major, L.H. News stories and images of immigration online: A quantitative analysis of digital-native and traditional news websites of different political orientations and social media engagement. *Atl. J. Commun.* **2023**, *31*, 207–226. [CrossRef]
19. Pournaki, A.; Willaert, T. Extracting narrative signals from public discourse: A network-based approach. *Humanit. Soc. Sci. Commun.* **2025**, *12*, 1774. [CrossRef]
20. Bodaghi, A.; Oliveira, J. Decoding semantic news ecosystem: Macro centrality analysis in Twitter news entity graphs and global influences. *Int. J. Digit. Libr.* **2025**, *26*, 18. [CrossRef]
21. Peivandizadeh, A.; Hatami, S.; Nakhjavani, A.; Khoshsima, L.; Qazani, M.R.C.; Haleem, M.; Alizadehsani, R. Stock market prediction with transductive long short-term memory and social media sentiment analysis. *IEEE Access* **2024**, *12*, 87110–87130. [CrossRef]
22. Liu, J.; Li, H.; Shen, W.; Sun, Y.; Yu, L.; Zhu, L.; Shi, Q.; Liu, Q. Predicting Online Group Opinions: A hypergraph-enhanced structure deep clustering with LSTM. *Comput. Hum. Behav.* **2025**, *172*, 108770. [CrossRef]
23. Bahari, A.; Zhang, X.; Ardasheva, Y. Establishing a nonlinear dynamic individual-centered language assessment model: A dynamic systems theory approach. *Interact. Learn. Environ.* **2023**, *31*, 4034–4056. [CrossRef]
24. Smith, C.B. The semantic attack surface: A systems-dynamic model of narrative in cyberspace. *IEEE Trans. Technol. Soc.* **2022**, *4*, 146–157. [CrossRef]
25. Parkinson, B. Heart to heart: A relation-alignment approach to emotion's social effects. *Emot. Rev.* **2021**, *13*, 78–89. [CrossRef]
26. Butler, E.A. Coordination in interpersonal systems. *Cogn. Emot.* **2022**, *36*, 1467–1478. [CrossRef] [PubMed]
27. Battista, D. Digital spaces and political antagonism: Towards a new ecosystem of dissent. *Soc. Issues* **2026**, *4*, 12–21.
28. Shah, D.; Ranka, H.; Ng, L.H.X.; Mehta, S. Bridging nodes and narrative flows: Identifying intervention targets for disinformation on Telegram. *arXiv* **2024**, arXiv:2411.05922. [CrossRef]