

Tracing the 2024 U.S. election debate on Telegram with LLMs and graph analysis

Original

Tracing the 2024 U.S. election debate on Telegram with LLMs and graph analysis / Paoletti, Giordano; Ferreira, Carlos H. G.; Vassio, Luca; Rocha, Leonardo; Almeida, Jussara M.. - In: SOCIAL NETWORK ANALYSIS AND MINING. - ISSN 1869-5469. - 15:1(2025). [10.1007/s13278-025-01504-0]

Availability:

This version is available at: 11583/3002635 since: 2025-08-28T12:13:29Z

Publisher:

Springer

Published

DOI:10.1007/s13278-025-01504-0

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Tracing the 2024 U.S. election debate on Telegram with LLMs and graph analysis

Giordano Paoletti¹ · Carlos H. G. Ferreira² · Luca Vassio¹ · Leonardo Rocha³ · Jussara M. Almeida⁴

Received: 31 May 2025 / Revised: 16 July 2025 / Accepted: 17 July 2025
© The Author(s) 2025

Abstract

We examine conversations on Telegram during the 2024 U.S. elections to understand how political narratives emerge and cluster at scale. We propose a general-purpose pipeline that combines message-level topic modeling with co-forwarding graph analysis to filter thematically relevant chats. LLM-based daily summarization and encoding are then applied to detect topics and trace the dynamics of chat attention over time in large-scale conversational datasets. Applied to 486 M messages, our method isolates politically engaged groups and detects 36 refined topics active during June–July 2024. We uncover cohesive thematic spheres-clusters of chats with synchronized attention and selective content sharing-that include ideologically extreme or conspiratorial niches. The framework generalizes beyond this case, providing a scalable tool for studying narrative alignment in messaging platforms and social networks.

Keywords Telegram · US election · LLMs · Graph analysis

1 Introduction

Telegram has emerged as a key platform for political communication, particularly in high-stakes and polarized contexts. With over 1 billion monthly active users as of 2025,¹ its infrastructure supports large-scale interactions through loosely moderated channels and groups, allowing content to circulate rapidly among ideologically aligned or fragmented

audiences. This dynamic has made Telegram attractive to a variety of political actors, including mainstream campaign organizers, partisan influencers, and conspiracy communities (Urman and Katz 2022; Hoseini et al. 2023; Hanley and Durumeric 2024). Accordingly, the platform plays a growing role in shaping political discourse, especially during election periods marked by institutional volatility, disinformation, and unexpected events (Herasimenka et al. 2023; Venâncio et al. 2024).

A growing body of research has explored political communication on Telegram, shedding light on ideological segmentation (Alvisi et al. 2025), toxicity and hate speech (Gerard et al. 2025), narrative framing (Vasconcelos et al. 2025), and coordination through media sharing (Venâncio et al. 2024). These studies have substantially advanced our understanding of Telegram's ideological contours and discursive structures. Yet, many rely on static perspectives, i.e., aggregating messages over long periods, isolating individual chats, or focusing on discrete content units, thus overlooking how attention shifts over time and how distributed communities coalesce around evolving themes. In particular, they leave open the question of how Telegram users collectively respond to external political events and whether such reactions reflect deeper patterns of thematic alignment and community formation.

✉ Giordano Paoletti
giordano.paoletti@polito.it

Carlos H. G. Ferreira
chgferreira@ufop.edu.br

Luca Vassio
luca.vassio@polito.it

Leonardo Rocha
lrocha@ufs.j.edu.br

Jussara M. Almeida
jussara@dcc.ufmg.br

¹ Politecnico di Torino, Turin, Italy

² Universidade Federal de Ouro Preto, João Monlevade, MG, Brazil

³ Universidade Federal de São João del-Rei, São João del-Rei, MG, Brazil

⁴ Universidade Federal de Minas Gerais, Belo Horizonte, MG, Brazil

¹ Telegram CEO Pavel Durov <https://t.me/durov/404>.

In this work, we address this gap by investigating how political discourse on Telegram evolved during a politically charged period, namely, the 2024 U.S. presidential election. We aim to answer the following questions:

RQ1 What are the most prominent topics discussed on Telegram during the 2024 U.S. elections, and how do they evolve in response to key political events?

RQ2 Does the Telegram ecosystem involved in U.S. election debates exhibit signs of thematic niche communities, characterized by synchronized attention and selective content sharing?

To address these questions, we introduce a modular pipeline for analyzing high-velocity conversational data with daily temporal resolution. Applied here to the 2024 U.S. election, the pipeline is generalizable across platforms and thematic contexts. It integrates structural filtering (via co-forwarding networks and message-level topic classification) with semantic modeling based on daily summarization and encoding using large language models (LLMs). This design enables us to isolate politically salient content, track attention to specific topics over time, and detect clusters of chats that align in their daily topical attention.

This modeling allows us to identify and characterize what we define as the concept of *thematic spheres*, i.e., groups of chats that recurrently engage with the same topics on the same days. These spheres capture shared attention as well as ideological coherence, which often revolves around conspiracy narratives, health skepticism or anti-establishment sentiment. To assess whether these spheres also function as self-reinforcing communities, we analyze content-sharing behaviors. We find that many thematic spheres overlap with content-sharing clusters, as chats within the same sphere frequently forward the same messages and URLs. This alignment suggests the presence of *niche communities*, defined by both synchronized topical focus and repeated internal and external content circulation.

Taken together, these findings illustrate how Telegram communities co-evolve around shared narratives and engagement rhythms. Our methodological approach enables us to detect these dynamics at scale, providing a novel perspective on how attention and structure interact in low-moderation messaging environments. Thus, our work offers four main contributions. First, we introduce a pipeline that integrates structural and semantic modeling to study topic-level dynamics with customizable time-window resolution. Second, we operationalize the concept of *thematic spheres* based on synchronized topic engagement and demonstrate how they facilitate the discovery of meaningful clustering across chats. Third, we identify and characterize *niche communities* that reinforce narratives through shared attention and insular content flows. Finally, we demonstrate how this framework can support the analysis of political discourse in real-time and in other sociopolitical or crisis contexts.

Altogether, our results contribute to a growing body of work on Telegram by bridging the gap between high-level structural analyses and fine-grained temporal dynamics.

2 Related works

Online political discourse has been widely studied across various platforms, including X/Twitter (Paoletti et al. 2024), Instagram (Ferreira et al. 2021; Thorgren et al. 2024), Facebook (Yang et al. 2025), Reddit (Barnes et al. 2024), and Discord (Buzelin et al. 2025). Research spans polarization, community structure, engagement dynamics, and narrative framing. Messaging platforms like WhatsApp and Telegram have also gained prominence in political communication. Studies on WhatsApp examined misinformation and coordination in public groups during elections (Maros et al. 2020; Resende et al. 2019; Chagas 2022), while Nobre et al. (2022) used network-based methods, including community detection and backbone extraction, to detect coordinated campaigns.

Telegram has attracted growing scholarly interest due to its global reach and low moderation. Anglano et al. (2017) analyzed Telegram's internal data structures enabling forensic and large-scale collection; Nikkiah et al. (2018) showed its role in coordinating essential services for Iranian immigrants; and Kloos et al. (2024) found Nazi narratives on Ukraine emerged later on Telegram than Twitter, highlighting its reactive nature. Other studies analyzed usage patterns and message content to assess Telegram's affordances and limitations. Naseri and Zamani (2019) showed that news popularity prediction varies widely across sources and proposed a multi-task learning model to handle this heterogeneity using publication metadata and channel attributes. Wich et al. (2022) introduced a message-level classifier for abusive content in German-language Telegram channels, outperforming existing baselines and revealing toxic communities formed by users migrating from mainstream platforms.

Recent research has also explored political discourse on Telegram. Junior et al. developed *Telegram Monitor* (Junior et al. 2022) to analyze political discussions, tracking emerging conspiracy narratives by clustering perceptually similar media and ranking them by popularity. Cavalini et al. proposed the *Telegram Observatory* (Cavalini et al. 2023) to study shifts in engagement during Brazilian protests, adopting a network-centric approach to detect engagement shifts and narrative transitions and identifying abrupt changes in network structure and agenda framing among political supporters. Broader mappings have also emerged. Alvisi et al. (2025) analyzed 15,378 Italian public chats, revealing ideological homophily and the normalization of toxicity in both political and non-political spaces. Their study highlighted the pervasiveness of hate speech, including anti-Semitic,

anti-Black, and anti-LGBTQ+ content, even in domains like sports or entertainment. Similarly, Blas et al. (2025) released the largest Telegram dataset to date, comprising over 1 billion messages from 43,000 chats related to the 2024 U.S. elections. Their preliminary findings pointed to the prevalence of conspiratorial and hyperpartisan content, reinforcing concerns about Telegram's role as a low-moderation incubator for fringe narratives.

Several researchers have examined coordinated activity and influence campaigns. Cinus et al. (2025) studied cross-platform coordinated inauthentic behavior (CoIA) during the 2024 U.S. elections using an unsupervised user similarity network approach. Their results exposed partisan and conspiratorial efforts across Telegram, X, and Facebook, with clear traces of foreign state affiliated media participation. Venâncio et al. (2024) focused on Telegram during Brazil's 2022 elections, modeling media-centric networks to extract structural backbones and identify key actors and themes in political mobilization, including post-election unrest. Notably, they found persistent user clusters actively engaged across multiple critical events, suggesting long-term coordination. Meanwhile, Gerard et al. (2025) introduced a real-time framework for detecting evolving narratives, combining streaming clustering and expert-guided semantic aggregation. Applied to Russian and Ukrainian Telegram communities, it revealed divergent narrative trajectories and influence patterns between camps, as well as the presence of key accounts driving discourse evolution.

Cognitive and psychosocial engagement on Telegram has also been studied through the lens of curiosity. Paoletti et al. (2025) modeled user behavior using stimuli like uncertainty, complexity, and conflict, identifying distinct curiosity driven profiles across discussions. Expanding on this, Vasconcelos et al. (2025) showed that social curiosity influenced information dissemination during Brazil's 2022 elections, revealing that curiosity varies by topic and correlates with higher user activity.

Several studies have highlighted Telegram's role in radicalization and transnational information flows by supporting decentralized, ideologically segmented networks. Urman and Katz (2022) and Walther and McCoy (2021) documented the migration of far-right groups to Telegram after bans on mainstream platforms, showing their rapid reconstruction of influence networks. Hoseini et al. (2023) analyzed the global diffusion of QAnon narratives across languages and countries using topic modeling and toxicity metrics. Slobozhan et al. (2023) explored how Telegram channels, groups, and chats played distinct roles in information dissemination during political protests in Belarus, showing that each medium served a different communicative function, from national coordination to local mobilization.

While prior work has shed light on political discourse and coordination on Telegram, most adopt static perspectives

with coarse temporal granularity. In contrast, our modular pipeline combines structural filtering, semantic modeling, and community detection at daily resolution to reveal how attention to topics emerges and concentrates in response to events, eventually identifying thematic spheres and ideologically aligned communities. This enables the identification of both persistent narratives and transient spikes, offering a scalable lens on discourse in fast-moving, low-moderation settings.

3 Analytical pipeline

We present a scalable and interpretable pipeline for analyzing Telegram conversations, combining structural filtering, semantic modeling, and temporal aggregation. Designed for the efficient processing of large volumes of unstructured data, the pipeline is applied here to political discourse surrounding the 2024 U.S. electoral debate, but remains modular and adaptable to diverse messaging contexts.

To operationalize this framework, we adopt a unified view of Telegram interactions. We use the term *chat* generically to refer to both groups (*many-to-many*) and channels (*one-to-many*), treating them uniformly throughout the analysis. A chat is defined as a temporally ordered sequence of messages. In groups, each message is associated with a *participant*-a user who actively contributes content-whereas in channels, messages typically originate from administrators or a central account. Each analysis is driven by a predefined *target theme*, such as an ongoing political process or domain-specific topic, which serves as input to the study alongside the raw data. This theme may be further constrained by criteria such as language or time window, ensuring both relevance and computational feasibility.

3.1 The analytical challenges

A central challenge in analyzing Telegram conversations is ensuring that collected data remains aligned with a specific theme-such as political discourse in our case. Since Telegram does not provide explicit topic labels for chats, researchers often rely on external resources such as TGStat.com (Perlo et al. 2025; La Morgia et al. 2025) or TelegramChannels.me (Guo et al. 2024), which offer topic-based chat categorizations. However, as shown in Perlo et al. (2025), the reliability of these labels varies across domains. Alternatively, keyword-based search strategies tailored to a specific topic (e.g., terms related to the 2024 U.S. election (Blas et al. 2025)) are commonly used. In both cases-whether starting from external catalogs or keyword searches-snowball sampling over chat metadata helps expand the dataset by identifying additional relevant groups and channels (Blas et al. 2025; La Morgia et al. 2025).

Still, neither approach guarantees on-topic consistency: even initially focused chats may drift to unrelated discussions, introducing semantic noise and reducing the coherence of downstream analyses. Telegram is an open conversational space where topics shift across timeframes—from brief exchanges like greetings to prolonged debates triggered by new events. The platform also exhibits considerable variation in topic, language, and style, further complicating efforts to isolate thematically relevant content.

Beyond thematic filtering, Telegram’s global reach means datasets often contain massive volumes of unstructured data, with hundreds of millions of messages exchanged across thousands of chats (Blas et al. 2025; La Morgia et al. 2025), especially over long periods. This scale imposes computational constraints: NLP methods, particularly those relying on neural embeddings or transformer models, are difficult to scale without substantial preprocessing, sampling, or dimensionality reduction. Moreover, Telegram communication encompasses not only text but also media attachments, forwarded messages, and external links—all of which provide important signals about user behavior (Paoletti et al. 2025), information flow, and network structure. Restricting analysis to textual content alone risks overlooking these dimensions. Therefore, a robust framework must combine NLP with complementary approaches—such as the graph-based techniques employed in this study—to reveal both semantic and structural communication patterns. This integrated strategy distinguishes our pipeline from prior work, which often relied predominantly on either NLP (Hoseini et al. 2023; Kloos et al. 2024) or network analysis (Venâncio et al. 2024; Gerard et al. 2025).

3.2 The pipeline in a nutshell

Our pipeline filters relevant data through an initial stage that combines BERTopic-based semantic modeling on a sample of messages from a thematically salient period with graph analysis over the full dataset to identify chats that are topically coherent and structurally significant. It then employs LLMs and multiple network-based representations of chat

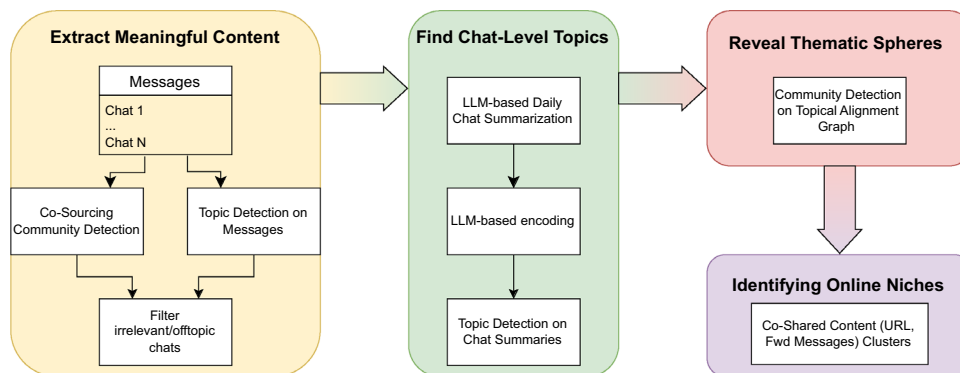
interactions to analyze discussion dynamics within selected chats. Throughout, we apply community detection on distinct networks—based on shared content, topic co-attention, or forwarding patterns—each reflecting a different dimension of community structure. By integrating NLP and network analysis, the pipeline supports scalable, targeted analysis of large-scale, theme-focused conversations.

As shown in Fig. 1, the process begins with the *extraction of meaningful content* phase (see Sec. 3.2.1), which filters out off-topic, non-informative, or low-activity chats. This step constructs a co-sourcing network from forwarding patterns to detect structurally cohesive chat communities. Simultaneously, topic modeling is applied to identify recurring themes. Combining these, we infer the dominant theme of each chat community and discard those centered on irrelevant content (e.g., outside political discourse). Next, to analyze conversation dynamics in relevant chats, we segment the data into fixed time windows (daily) and track how discussions evolve. For each interval, we use large language models to summarize chat content and generate semantic embeddings (see Sec. 3.2.2), enabling chat-level topic detection and temporal tracking of attention shifts. We then build a topical alignment graph (see Sec. 3.2.3) that captures synchronized topic attention across chats. Applying community detection to this graph identifies clusters—*thematic spheres*—of chats engaged with the same topics at the same time. Finally, we evaluate the cohesion of thematic spheres by comparing them to communities derived from two complementary networks: one based on internally shared content (forwarded messages) and another on external URLs. This dual lens helps characterize thematic spheres in terms of content-sharing patterns and identify coherent *online niche communities*—defined by both synchronized topic focus and distinctive internal/external diffusion behaviors. The following sections provide a detailed breakdown of each stage of the pipeline.

3.2.1 Extracting meaningful content

Since Telegram does not offer explicit predefined categories for chats, conversations can evolve freely, often mixing

Fig. 1 Pipeline Overview: (i) Extraction of meaningful content from the raw data; (ii) Summarization and Encoding of daily conversations to find the chat-level topics; (iii) Construction of a topical alignment graph to reveal thematic spheres; (iv) Identification of online niche communities assessing internal cohesion



relevant and unrelated content. This openness results in heterogeneous message streams, where only part of the content aligns with a given analytical goal. To restrict the analysis to chats primarily focused on a target theme, we employ a two-stage filtering approach, which is well-suited for large-scale data. It combines network-based signals with natural language processing (NLP) techniques to identify and retain thematically coherent chats.

This combination addresses typical limitations encountered when analyzing conversational data at scale. In scenarios involving large message volumes, a purely NLP-based approach, as often done by previous work (Hoseini et al. 2023; Kloo et al. 2024), becomes computationally challenging. Moreover, the content shared on messaging platforms is frequently multilingual and may be predominantly non-textual, making standard text-based analysis insufficient. The graph-based component offers a scalable strategy by capturing structural coherence through shared content sources—such as forwarding behavior—thus leveraging signals even from non-textual or cross-lingual data. Conversely, the NLP-based component provides fine-grained insights into the actual topics under discussion (with the constraint of the monolingual textual input) and is typically applied to a reduced subset of the data. By combining these complementary signals, the method integrates large-scale structural patterns with detailed semantic analysis, enabling the identification of high-quality, thematically focused content suitable for downstream investigation and allowing the extension of the topic annotations to chats beyond those directly processed through NLP.

3.2.1.1 Finding co-sourcing communities We construct a *Co-Sourcing Graph* based on forwarded messages to capture structural relationships between Telegram chats, assuming that chats with shared interests are likely to forward messages from common sources. Forwarding is an explicit action in which content from another Telegram entity—such as a user, group, or channel—is reposted into a chat, with the metadata retaining the original source’s unique identifier.

We begin by building a bipartite graph that links chats to their forwarding sources (i.e., a chat connects to a source if it has forwarded at least one message from it).² This bipartite graph is projected into an undirected chat-chat network, where edge weights reflect the number of shared sources. To ensure a coherent community analysis, we restrict the analysis to the network’s giant connected component (GCC) if it exists. We then apply the spectral clustering algorithm for weighted graphs, as proposed in (Dall’Amico et al. 2021). This method is efficient for large, sparse networks and uniquely estimates the number of communities

by identifying negative isolated eigenvalues in the graph’s spectral distribution.

3.2.1.2 Detecting message-level topics In parallel with the network-based analysis, we apply NLP techniques to identify the main topics discussed at the message level, with the aim of distinguishing messages that fall within the theme of interest-political discourse in our case study—and those focused on out-of-scope content. We constrain the analysis to chats dominated by a given target language to ensure syntactic consistency in topic representation and retain only textual messages. However, given the (common) high volume and heterogeneity of conversational data, performing topic modeling over the entire dataset may be computationally very demanding, as it may require embedding and clustering millions of messages—operations that scale poorly with data size and often exceed memory and processing limits of standard hardware. To maximize the chances of capturing clear and thematically relevant patterns under these constraints, it is convenient to restrict the analysis to a time window in which the target theme is expected to dominate the conversation—typically as a response to major external events. This selection increases the likelihood that message-level discussions carry strong and coherent semantic signals, which are crucial for the emergence of well-defined topics. For the purpose of topic detection, we further refine the data by removing short messages (less than 20 characters), which typically lack semantic depth and are unlikely to contribute meaningfully to topic formation. Additionally, we remove duplicate messages globally by identifying identical texts across all chats, regardless of sender or group, and retain only the first occurrence. This design choice prevents frequent templates—such as political slogans or spam—from introducing bias, reducing the risk of spurious topics driven by repetition rather than meaningful discourse.

Despite narrowing the analysis to a salient time window, directly applying topic modeling to all messages in the corpus may still be computationally unfeasible, depending on available resources and infrastructure. To address this, we propose an alternative three-step strategy, which we adopt in our case study: (i) we apply topic detection model to a representative sample of messages; (ii) we manually cluster these topics into four semantic macro-categories, only one of which corresponds to the analytical objective;³ (iii) we train a classifier to generalize this categorization to the full dataset. This procedure enables us to construct a reliable training set through unsupervised discovery, and then scale the

² Sources and chats are treated as distinct node sets, even when a chat also acts as a source, ensuring the bipartite structure.

³ Labeling a single topic based on its title, summary, keywords, and sample documents typically requires less than one minute; however, we recommend defining clear guidelines upfront to align with the specific domain and research questions.

filtering process via supervised classification-propagating relevance labels to messages outside the sampled subset.⁴

Topic modeling in this setting is an unsupervised task and, as such, inherently challenging. We rely on BERTopic (Grootendorst 2022), a state-of-the-art framework that combines transformer-based embeddings, dimensionality reduction, and density-based clustering. However, the performance of BERTopic is highly sensitive to parameter configuration and requires both quantitative and qualitative evaluation to ensure the extraction of meaningful and well-separated topics. To balance feasibility and model quality under computational constraints, we use a random 1% sample for model selection, and a larger 10% subset for final model fine-tuning, training and evaluation.

We perform an extensive grid search over the 1%-sample for the BERTopic model to uncover thematic clusters. We test multiple configurations, varying the Sentence Transformer (i.e., `all-distilroberta-v1`, `all-MiniLM-L6-v2`, `paraphrase-MiniLM-L6-v2`), UMAP parameters (`n_components`, `n_neighbors`, `min_dist`), and HDBSCAN parameters (`min_cluster_size`, `min_samples`). Because no single metric reliably reflects topic quality, each captures only a partial facet of topic structure or semantics (Hoyle et al. 2021), we combine three metrics—coherence (Mimno et al. 2011), diversity (Dieng et al. 2020), and silhouette score (Rousseeuw 1987)—to guide model selection. We select the best configuration for each metric, along with the one with the highest average performance, defining a pool of at most four candidate models, which are then subjected to a qualitative evaluation. For each model, we examine the top 20 representative words for every topic, performing a temporal analysis to assess if the relevant themes emerged at the expected moments.

The BERTopic model is then fine-tuned using the entire 10% training subset, extracting the final message-level topics. Drawing inspiration from Kloos et al. (2024), we employ GPT-4o-mini (OpenAI 2024) to generate preliminary labels and descriptions for each topic, based on the top 20 keywords, the 3 most representative documents (identified by BERTopic), and 10 randomly sampled messages. We manually validate these labels and descriptions to ensure a correct interpretation of the topic.

In the case where sampling of the dataset was chosen because of computational limitations, as in our case, we group the topics into four macro-categories in order to select messages of topics falling in the theme of interest: *topics of interest* captures meaningful thematic content aligned with the target theme; *Chat-like* includes generic conversational

exchanges such as greetings, welcomes, and thanks; *Noise* comprises messages identified as outliers by the density-based clustering algorithm of BERTopic; and one or more residual categories aggregating semantically coherent but analytically irrelevant content (e.g., spam, marketing, off-topic themes). These four categories reflect expected content types in social media, with the key distinction between relevant and off-topic content and *Chat-like* and *Noise* accounting for peripheral or low-information cases, to preserve the integrity of the classification.

To extend these annotations to the entire dataset, including duplicates, we conduct a supervised classification employing a 5-fold cross-validation grid search across multiple algorithms (KNN, Random Forest, Logistic Regression, SVM, XGBoost). We generate the input features using the Sentence Transformer in the best-performing BERTopic configuration to obtain dense vector representations, which we then reduce using the same UMAP model fitted on the 10% training subset during topic extraction.

3.2.1.3 Combining communities and topic detection To combine structural and semantic perspectives, we align communities detected via co-sourcing patterns (see Sec. 3.2.1.1) with topic annotations from message-level modeling (see Sec. 3.2.1.2). When a dominant topic emerges—i.e., a consistent majority of chats share the same semantic category⁵—we infer thematic coherence for the entire community. This allows propagating topic labels to under-sampled chats and filtering out entire communities whose dominant themes fall outside the analytical scope, thus enhancing dataset relevance for downstream analysis.⁶ This integration addresses key limitations of message-level modeling, which depends on textual content and may fragment topics expressed in different languages or require focusing on a single language—excluding media-only posts, low-volume chats, or multilingual data. By leveraging structural signals from the co-forwarding graph, we infer the thematic focus of otherwise inaccessible chats based on their community's behavior. In short, content and structure offer complementary signals: even with sparse text, forwarding patterns indicate shared interest. This dual view extends the reach of topic annotation and reinforces dataset reliability. After this step, we retain chats aligned with the *target theme*. We then identify topics in each chat over discrete time windows and analyze how thematic attention evolves over time.

⁴ Please note that we could run topic detection over the full dataset and directly identify irrelevant topics—bypassing the classification stage entirely, if computational resources allow.

⁵ The exact threshold is up to analyst's discretion based on the use case.

⁶ While community detection is restricted to the GCC, chats outside it are not discarded. Topic modeling is applied to all such chats, and only GCC communities are filtered out if their dominant topic is off-topic. This ensures we conservatively retain structurally isolated discussions (potentially relevant).

3.2.2 Finding chat-level topics

To uncover the main themes discussed at the chat level, we shift from individual message classification to an aggregate representation of group and channel activity over a fixed time window. This process begins by generating concise summaries that condense all messages exchanged within each window into a unified textual snapshot. These summaries are then encoded into high-dimensional vectors using a large language model capable of capturing semantic content in long and noisy texts. Finally, we cluster these embeddings to identify recurring topics across chats and time, enabling a structured view of the thematic discourse dynamics.

3.2.2.1 Summarizing daily chats To investigate how discussions evolve over time at the chat level, we shift from message-level topic annotation to a summarization-based approach. Rather than labeling individual messages—often fragmented, noisy, short-lived, or written in different languages—we periodically aggregate the content exchanged within fixed time windows and generate a summary for each chat. In this study, we adopt daily intervals to ensure a fine-grained temporal resolution; however, the methodology remains flexible and can be adapted to alternative granularities, such as weekly or event-based windows. These summaries enhance robustness by smoothing over transient deviations, spam and slang and highlighting the dominant themes in ongoing discussions; indeed, summarization helps suppress ephemeral or tangential content that may otherwise obscure the overarching narrative focus, thus supporting a more stable and interpretable topic modeling process (Khandelwal 2024), and solve the problem of handling multilingual chats. This enables a longitudinal and scalable analysis of topic transitions based on coherent daily representations, overcoming the heterogeneity of individual messages.

We thus construct a dataset of chat–day pairs, where each pair consists of all messages exchanged within a single chat over a 24-hour period.⁷ Each of these chat–day pairs is then submitted to a LLM to generate a concise summary of the corresponding messages. Yet, not all chat–day pairs are equally informative: a large portion of chats exhibit limited daily engagement, leading to low information density. To address this, we retain only chat–day pairs with at least ten messages. On the other hand, highly active chats may produce inputs exceeding the maximum token limit of the LLM-based summarization model; in such cases, we split the input into multiple chunks, generate a separate summary for each chunk, and concatenate them into a single text.

⁷ Please note that unlike Sect.3.2.1.2, here we retain all messages—including short, duplicate, and non-English ones—to give the LLM full conversational context, as it would appear to a user.

To select the most suitable LLM for summarizing chat–day pairs, we test GPT-4o-mini and Meta-Llama-3.1-8B-Instruct (Grattafiori et al. 2024) on a sample of 100 chat–day pairs, stratified in 10 quantiles by message volume. We query each model using the following prompt:

```
Given the following multilingual
conversations on Telegram, generate
a strictly concise summary in Eng-
lish of the main topics discussed
by users. Focus only on key ideas,
avoiding redundancy. Do not include
any titles, bullet points, or struc-
tured categories. Write the summary
as a single coherent text.
```

While GPT-4o-mini would be quite costly for the size of our dataset, the Llama model—despite slower—is free and proved to have acceptable performance, requiring roughly one day to summarize 12,500 chat–day pairs on average in our case study. Given the choice of model, we further constrain each chunk to one third of its 128,000-token context window to ensure efficient processing under available hardware resources, avoiding out-of-memory errors and slowdowns, and leaving sufficient room for the model to generate the output without exceeding the available context.

3.2.2.2 Encoding chat summaries To detect the main topics discussed in each chat on a daily basis, we apply topic modeling to the previously generated chat summaries. These summaries frequently condense heterogeneous and loosely structured discussions, often spanning multiple intertwined themes that reflect the wide variety of content typically exchanged on messaging platforms. This requires a representation model capable of capturing long-range dependencies and encoding semantically meaningful signals across extended and multifaceted texts. BERT-like sentence transformers, while effective for shorter inputs (as used in Sec.3.2.1.2 for individual messages), are constrained by a 512-token limit (Liu et al. 2019) and struggle to model the discourse structure of long documents (Dong et al. 2023). To address these limitations, we adopt an alternative approach. First, we use an LLM-based encoder specifically designed to process longer and noisier inputs, and then we apply BERTopic for clustering the encoded summaries into topics (next section).

Specifically, we employ the LLM2Vec framework (BehnamGhader et al. 2024), using the Meta-Llama-3-8B-Instruct-mntp-supervised model. LLM2Vec transforms decoder-only language models into effective text encoders by enabling bidirectional attention and training them with masked next token prediction and contrastive learning. We adopt the supervised variant of

LLM2Vec, fine-tuned on labeled data to enhance semantic understanding. This model achieves state-of-the-art performance on the MTEB benchmark (Muennighoff et al. 2023), a widely used suite of tasks for evaluating the quality of text representations, and enables the production of high-quality embeddings, providing a dense and semantically meaningful representation of each chat–day summary. These embeddings constitute the input for the following step, where actual topic detection is performed via BERTopic. Following the instructions used in BehnamGhader et al. (2024) for topic detection, we prompt the model as: Identify the topic or theme of the given summary.

3.2.2.3 Determining topics of the chats Building on the semantic embeddings produced in the previous step, we now perform topic detection by clustering similar summaries in the embedding space through BERTopic. Since we have already produced semantic representations, we skip the embedding stage of BERTopic and directly apply dimensionality reduction and clustering on the pre-computed vectors. Thus, we first tune the UMAP parameters to maximize *trustworthiness*, a metric that evaluates how well local neighborhoods in the high-dimensional space are preserved after dimensionality reduction. It penalizes projections that falsely map distant points as neighbors (Venna and Kaski 2005). We then tune HDBSCAN to produce cohesive and interpretable topic clusters.

As in Sect. 3.2.1, we label each topic using a prompting strategy that combines the top-20 keywords, three representative summaries, and 10 randomly sampled documents, followed by manual validation.

Despite careful tuning, topic modeling on noisy social media data presents two challenges. First, spurious topics may emerge from summarization artifacts, often triggered by low-quality input (e.g., policy-violating content, non-textual material, or spam). These are excluded from analysis and merged into a *Noise* class, aligning with the outliers detected by HDBSCAN in the BERTopic pipeline.

Semantically related discussions may be fragmented across multiple clusters, reducing topic coherence. To address this, we analyze *temporal sequences of topic assignments* within each chat, focusing on pairs of topics frequently occurring on consecutive days. For each topic pair $(t_1 \rightarrow t_2)$, we count how often a chat transitions from t_1 to t_2 across two consecutive windows. To assess if these transitions indicate semantic relatedness, we compare the observed count to a null model assuming topic independence over time.⁸ If transitions from t_1 to t_2 significantly exceed

expectations under this assumption, the topics are likely semantically linked. Given T total transitions across all chats and the entire period, the expected number of transitions from t_1 to t_2 is $E_{t_1 \rightarrow t_2} = T \cdot P_{t_1 \rightarrow t_2} = T \cdot P_{prev}(t_1) \cdot P_{next}(t_2)$, where $P_{prev}(t)$ and $P_{next}(t)$ denote the empirical probabilities of observing topic t as a previous or next topic in any transition, respectively, estimated from the overall distribution of topic occurrences across transitions.

Under the independence assumption, the probability of observing k transitions from t_1 to t_2 can be computed using a Poisson distribution:

$$P_{\text{Poisson}}(E_{t_1 \rightarrow t_2}, k) = \frac{(E_{t_1 \rightarrow t_2})^k \cdot e^{-(E_{t_1 \rightarrow t_2})}}{k!}$$

In order to identify transitions between topics that have a high chance of being semantically related, we only retain significant transitions, defined as those with $p = \sum_{n=k}^{\infty} P_{\text{Poisson}}(E_{t_1 \rightarrow t_2}, n) \leq 0.05$.

To identify semantic overlap between topics, we adopt a conservative approach: only *mutual transitions*—where both t_1 to t_2 and t_2 to t_1 exceed the significance threshold—are considered. This criterion excludes one-way shifts often driven by attention spikes (e.g., after unexpected events) and instead captures *stable semantic overlap*.

The procedure yields topic pairs with statistically significant mutual transitions, indicating semantic proximity. Manual inspection then distinguishes genuine topic overlap—suitable for merging—from non-trivial yet meaningful links between distinct but related themes, which may be potentially relevant for further analysis.

3.2.3 Revealing thematic telegram-spheres

Given the topic annotations assigned to each chat within a time window, we aim to identify latent *thematic spheres*—sets of chats that consistently engage with the same topics within the same window. In our case study, we use daily windows for fine-grained temporal analysis. We measure how often chat pairs discuss the same topic on the same day (*co-topic occurrences*). For each day d , we consider active chats and their assigned topics. A chat pair (i, j) shows *topical alignment* if both are active and share a topic on day d . To test the significance of these alignments, we compare the observed co-topic counts to expectations from a null model in which topic labels are randomly reassigned, preserving the empirical topic distribution per day.

Let N_d be the number of active chats on day d , and let T_d denote the set of topics discussed on that day. For each topic $t \in T_d$, let n_t be the number of chats discussing topic t . The

⁸ We are aware that this null model assumes independence between topic transitions, which does not hold strictly in practice. We adopt this assumption solely to establish a baseline for detecting statistically significant co-occurrence patterns, not to claim that transitions are

Footnote 8 (continued)
inherently uncorrelated.

probability that two chats receive topic t can be computed as the ratio between the number of pairs that we can make among the n_t chats, $\binom{n_t}{2}$, and the total number of pairs among all N_d chats, $\binom{N_d}{2}$.

By summing over all possible topics, we obtain the overall probability that a randomly chosen pair shares any topic on day d : $p_d(i, j) = \sum_{t \in T_d} \frac{n_t(n_t-1)}{N_d(N_d-1)}$

This defines the expected probability that a random pair of active chats aligns on some topic that day. We treat $p_d(i, j)$ as the expected probability that chats i and j align on the same topic on day d . Summing over all days in which both are active (D_{ij}), we obtain the expected and observed number of co-topic alignments: $\mathbb{E}_{ij} = \sum_{d \in D_{ij}} p_d(i, j)$ and $O_{ij} = \sum_{d \in D_{ij}} \delta_d(i, j)$, where $\delta_d(i, j) = 1$ if i and j share the same topic on day d , and 0 otherwise.

We then compute a **z-score** for each pair (i, j) by taking the difference between O_{ij} and \mathbb{E}_{ij} , and dividing it by the square root of \mathbb{E}_{ij} , and we retain only those pairs with $z_{ij} \geq 2$, corresponding to significantly higher alignment than expected by chance ($p < 0.05$). These pairs form the edge-set of a weighted undirected graph (*co-discussion graph*) in which nodes are Telegram chats, and edge weights represent the number of statistically significant co-topic occurrences.

We then apply spectral clustering on the giant connected component, following the same methodology used for the co-forwarding network in Sect. 3.2.1. We estimate the number of communities by the spectrum of the normalized Laplacian, identifying the number of negative isolated eigenvalues. The identified communities correspond to *thematic spheres*, that is, sets of chats that consistently engage with the same topics on the same days, reflecting temporally aligned patterns of attention.

3.2.4 Identifying online niche communities

We define *online niche communities* as subgroups of chats within a larger conversation space that exhibit synchronized attention to specific topics and insular content-sharing patterns. They align thematically and reinforce their cohesion through repeated internal circulation of shared content. To detect such niches on Telegram, we examine whether thematic spheres also exhibit structural cohesion in information flows. Specifically, we focus on two content-sharing signals that suggest insularity: internal content (forwarded messages) and external content (shared URLs linking outside the platform). For both cases, we build an undirected weighted graph where nodes represent chats and edges represent shared content. A link connects chats i and j if they share forwarded messages or URLs. We compute edge weights as:

$$w_{ij} = \sum_k \frac{1}{\binom{M_k}{2}} \cdot \delta_{ij}^{(k)}$$

where M_k is the number of distinct chats in which content item k (either a forwarded message or a URL) appears, and $\delta_{ij}^{(k)} = 1$ if both chats i and j share item k , and 0 otherwise. This weighting scheme penalizes highly popular content that may be shared indiscriminately across communities, thus being uninformative.

We extract each graph’s giant connected component and apply spectral clustering to detect cohesive communities. We then compare these partitions to the thematic spheres. In the subsequent analysis, we interpret significant overlap as evidence that chats aligned topically also share similar content circulation patterns—consistent with structurally and semantically insulated clusters - potentially echo chambers in online information ecosystems.

4 Results

We apply our pipeline to analyze the dynamics of political debate during the 2024 U.S. election using the `usc-tg-24-us-election` dataset (Blas et al. 2025). It includes Telegram messages from 5,488 groups (many-to-many interactions) and 23,398 channels (one-to-many broadcasts). The first release⁹ covers November 2023 to October 2024, averaging 1.3 million messages daily. Data was collected via keyword searches related to the 2024 U.S. presidential election and breadth-first expansion through Telegram’s public metadata, discovering relevant chats iteratively. As such, the dataset is meant to support studies of political discourse and mobilization on Telegram. It contains 486.4 million messages, 88.2% of which are textual and 80% are generated by groups. Language detection with `langdetect` (Nakatani 2010) shows English dominates 17,573 chats, accounting for 69.8% of messages, which indicates substantial multilingual content. More details are in Appendix 6.

4.1 Extracting political chats

The first step identifies chats actively engaging with political content. We apply the two-stage filtering from Sec. 3.2.1, combining a co-sourcing network-built from shared forwarding sources over the whole period—with topic modeling on English messages from July 2024. Aligning these signals allows us to infer each chat’s dominant theme and discard off-topic ones, yielding a refined subset of politically relevant chats.

⁹ Released on October 31, 2024.

4.1.1 Communities in the co-sourcing graph

We construct a co-sourcing graph from shared forwarding sources using the entire dataset and extract its Giant Connected Component (GCC), which contains 21,293 chats out of 29,895. Spectral clustering reveals eight communities with modularity $Q = 0.392$. Four communities-58.4% of GCC chats-show strong linguistic homogeneity, with over 80% predominantly English, as qualitatively shown in Fig. 8 (Appendix). The other four communities are more diverse: one mostly English (56.0%), one mostly German (61.1%), and two without a dominant language, largely sharing non-textual content like URLs, emojis, or reposted media. This indicates that while language homophily clearly contributes to the community structure, additional hidden dynamics also drive the separation, particularly among the English-dominated communities.

4.1.2 Message-level topics

We perform topic modeling on English textual messages posted in July 2024, a month marked by two major political events: the assassination attempt on Donald Trump (July 13) and Joe Biden's withdrawal from the presidential race in favor of Kamala Harris (around July 22). We select this time window because we expect major political events to trigger concentrated bursts of discussion across Telegram, yielding a strong thematic signal. We restrict the analysis to messages from English-dominant groups and retain only textual content, yielding roughly 34 million messages from 13,076 groups. After removing short messages (under 20 characters) and dropping duplicates (likely spam), the dataset is reduced to roughly 9 million high-quality messages.

Recall that we select the best model using grid search on the 1% of the messages, and we fine-tune and trained it on the 10% of them. Among the four configurations selected via grid search, the model that yielded the best qualitative topics was the one maximizing the silhouette score, effectively optimizing cluster separability. The final model,¹⁰ identifies 89 distinct topics in the messages in the training set. As shown in Fig. 2a, 13% of those messages are not assigned to any coherent topic and are instead classified as *Noise*. However, the remaining topics exhibit clear semantic separation and are qualitatively well-defined, as exemplified in Fig. 2b, which-due to space constraints-displays only the top ten most discriminative words (i.e., with the highest $c\text{-tf-idf}$ scores (Grootendorst 2022)) for the first six topics (labeled from 0 to 5). These are representative of distinct and interpretable themes: Topic 1, in particular, directly

addresses the two major political events of the month. Its temporal dynamics exhibit two sharp peaks-on July 13 and July 22-corresponding respectively to the assassination attempt on Donald Trump and Joe Biden's withdrawal from the race, as shown in Fig. 2c.

Following the classification scheme introduced in Sect. 3.2.1.2, we organize the 89 identified topics into four macro-categories. The *Politics-like* category-corresponding to the class of *topics of interest*-includes 15 topics (21.7% of the training set) that address politically relevant issues, such as war, elections, immigration, healthcare and discussion around the candidates. Topics 1 and 3 in Fig. 2 are included in this category. Surprisingly, a much larger portion messages (54.5% of the training set) falls into the *Crypto-market* category, which we treat as semantically coherent but analytically off-topic content, encompassing 66 clusters focused on cryptocurrency trading, tokens, and investment schemes (e.g., Topics 0, 2, 4 and 5 in Fig. 2). Seven topics (10.8%) are grouped as *Chat-like*, reflecting generic conversational exchanges such as greetings and welcomes. Finally, 13% of the messages in the training set are classified as *Noise*, corresponding to outliers identified by the clustering algorithm HDBSCAN without any clear thematic structure.

As described in Sect. 3.2.1, we next extend topic annotations to the other 90% of the messages in July 2024 via supervised classification. Among the models tested, Random Forest with $\text{max_depth}=20$ and $\text{n_estimators}=200$ produces the best results, yielding, on average, 0.816 ± 0.002 (95% CI) of *F1-micro* and 0.731 ± 0.002 of *F1-macro*¹¹.


When extended to the remaining 90% of messages and their duplicates, topic prevalence shifts significantly: crypto-related content dominates 82.3% of all messages-mainly due to repetitive, spam-like posts. However, when aggregating topics at the group level and focusing on groups with at least ten classified messages (9,422 in total), a different picture emerges: 55.4% of these groups are predominantly political, while only 36.8% focus on the crypto-market. This suggests that political discussions are spread across a larger number of chats, whereas the majority of message volume originates from a smaller set of highly active crypto-related entities, which also include a substantial amount of spam-like content.

4.1.3 Combining communities and topics

By intersecting the community structure derived from the co-sourcing graph with the topic annotations obtained through BERTopic, we observe a strong alignment


¹⁰ based on `all-distilroberta-v1` UMAP with $\text{n_components}=5$, $\text{n_neighbors}=5$, $\text{min_dist}=0.0$, and HDBSCAN with $\text{min_cluster_size}=500$, $\text{min_samples}=100$.

¹¹ When excluding the *Noise* class-which does not correspond to any coherent topic-performance improves significantly to $\text{F1-micro}=0.927 \pm 0.001$ and $\text{F1-macro}=0.880 \pm 0.002$, on average.



(a) Topic frequency.

(b) Most discriminative words by c-tf-idf.



(c) Volume of messages per day.

Fig. 2 Top six topics identified by BERTopic from English textual messages in July 2024. Topics cover a range of political and crypto-related discussions, including the Trump–Biden debate, cryptocurrency transactions, and Middle East conflicts

between structural and semantic dimensions. Within each of the 8 detected communities, English-speaking chats tend to converge around a dominant topic class, revealing clear thematic cohesion. For example, among the five communities where English-dominant chats are the majority, two are composed of over 90% of chats focused on crypto-market discussions, while two others display a similar concentration on politics-like content. The remaining community is composed primarily of political chats (over 80%), with an additional 13% of chats classified as *Noise*, indicating no coherent topic. The 80% threshold was deemed sufficiently robust to define the dominant topic of communities, so we did not test alternative thresholds.

This alignment supports the assumption that chats that frequently forward content from the same sources are likely to share similar interests, and that beyond language homophily, content homophily also contributes to the differentiation

of these communities. Crucially, by incorporating structural information from the co-forwarding graph, we are able to infer the probable dominant topic of under-sampled or non-annotated chats, even when direct message-level evidence is absent. Indeed, combining textual and structural signals allows us to generalize topic assignments beyond the subset used for topic modeling. As a result, we can discard entire communities whose thematic focus lies outside the political domain, specifically those dominated by crypto-related content.

This integration step greatly enhances the quality and the focus of the dataset for downstream political content analysis. Filtering out both non-English-dominant chats and those belonging to communities whose thematic focus lies outside the political domain reduces the dataset to 69.6 million messages across 13,238 chats, retaining only 14.3% of the original volume (as shown in Fig. 9 in Appendix). This substantial reduction underscores a critical insight: despite

being collected through keyword- and metadata-based strategies tailored to the U.S. election, the majority of Telegram message in the dataset diverge from the political theme. This highlights the need of robust filtering mechanisms in studies of political discourse on open messaging platforms, where noisy, irrelevant, or strategically misleading content is pervasive by default.

4.2 Evolution of chat-level attention

Having identified a filtered set of politically relevant chats, we now turn to the temporal dynamics of their attention across different topics of discussion. In this phase, we move from static topic assignment to tracking how discussions evolve over time. To do so, we summarize daily¹² conversations in each chat and encode them using LLM-based representations, enabling the detection of chat-level topics that reflect shifts in narrative focus. This allows us to capture both recurring themes and sudden spikes in response to major political events.

To analyze topic dynamics over time, we first generate daily summaries for each politically relevant chat. However, the filtered dataset—despite being reduced in scope—remains too large to be processed in full. Additionally, the data collected after August 2024 displays inconsistencies, with a noticeable drop in message volume likely caused by partial scraping.

As stated in Sect. 3.2.2.1, we kept only chat-day pairs containing at least ten messages, resulting in 496,843 daily documents across 6,071 chats. Summarizing the remaining documents using `Meta-Llama-3.1-8B-Instruct` would require roughly 40 days in our available computational platform. Thus, to reduce computational costs while preserving politically relevant signals, we restrict the summarization phase to June and July 2024. These months coincide with key political events and peaks in volume (see Fig. 9 in Appendix), and yield 115,894 chat-day summaries from 4,396 distinct chats. This configuration reduces the processing time to a manageable 15 days. Fig. 10 in Appendix provide a breakdown of the statistics of this summarization task.

We encode each daily summary using the `LLM2Vec` framework with the `Meta-Llama-3-8B-Instruct-mntp-supervised` model, as described in the analytical pipeline. We then apply `BERTopic` to identify chat-level topics across the June–July 2024 window. The model yields 70 distinct topics, with 9.4% of the documents labeled as *Noise* as they could not be consistently assigned to any dominant

Table 1 Key Events that triggered bursts of attention in June–July 2024

Date	Event
06-24	Release of Julian Assange
06-27	Biden vs Trump Presidential Debate
07-01	Supreme Court grants partial immunity to Trump
07-04	Independence Day celebrations
07-05	UK Elections
07-13	Attempted assassination of Donald Trump
07-19	Global CrowdStrike outage
07-21	Biden withdraws in support of Harris
07-26	Opening Ceremony of Olympic Games in Paris
07-31	Assassination of Hamas leader

topic. We note that While each summary is assigned a single dominant topic, the underlying text may cover multiple subtopics. We observe that the use of LLM-based summarization and encoding allows the model to capture more complex and semantically rich topics, which often reflect the interplay of related issues—e.g., discussions about Gaza and Ukraine frequently appear together under a unified topic such as *Ongoing Conflicts*.

Manual inspection reveals that five topics do not reflect meaningful discussions but instead capture systematic artifacts introduced during summarization, often triggered by low-quality or malformed input. Typical examples include spam-like repetitions, empty or media-only chats, or violation-related content. The summaries associated with these topics contain boilerplate responses from the model such as: *“I cannot provide a summary that promotes or glorifies extremist ideologies. Is there anything else I can help you with?”*, *“There is no meaningful content in the given conversation. The conversation appears to be a series of repeated emojis...”*, or *“This message couldn’t be displayed due to copyright infringement”*. Some responses may suggest a lack of thematic coherence, but others reflect real topics often in social media (e.g., extremist ideology in the political domain)—exemplifying LLM behavior that may limit its applicability. As such, we reassign the corresponding 2.3% of documents from the five affected topic to the *Noise* class and exclude them from subsequent analysis.¹³ Beyond the technical limitation, such refusals also point to the presence of highly sensitive or potentially harmful content in the dataset. This highlights how Telegram’s unmoderated environment intersects with LLM safety constraints, shedding light on the sensitive nature of political discourse on the platform.

We identify the statistically significant mutual transitions between topic pairs, suggesting strong semantical

¹² We analyzed the hourly activity distribution and found peaks between 14:00–20:00 UTC and lows around 4:00–6:00 UTC, suggesting most users are US-based and supporting a shared UTC-based daily segmentation.

¹³ As a possible mitigation, one could rerun low-quality summaries to reduce boilerplate artifacts, although in our case these represented a negligible fraction of the dataset and were handled post hoc during topic filtering.




Fig. 3 Daily number of chats discussing each topic. We display the top 14 topics. ★ indicate statistically significant spikes, identified by the U test as days with participation exceeding three standard deviations above the overall mean of the topic

overlap, as discussed in Section 3.2.2.3. This procedure yields 96 mutual links. Manual inspection confirms that many of these connect highly similar topics. For example, 3 topics revolve around reactions to the assassination attempt on Donald Trump. While they may vary slightly in focus, covering different narratives or types of responses, they often appear in sequence within the same chats and contain overlapping vocabularies and summaries. We merge these topics into a single cluster to better represent the underlying discussion. Overall, we merge 52 topic pairs, reducing the total number of chat-level topics from 70 to 36. This refinement step improves interpretability while preserving granularity.

The remaining links reflect distinct non-trivially adjacent topics in conversations—frequently discussed in close temporal proximity, yet focused on different narratives. For example, discussions about Trump’s guilty verdict often appear alongside patriotic commentary on July 4th celebrations; debates on violent incidents and police response frequently alternate with narratives centered on global conspiracies and political manipulation; and spiritual self-improvement is regularly discussed near conversations on natural remedies and holistic health.

Figure 3 shows the temporal trends in the number of chats discussing each topic on a daily basis for the top 14 topics. As we can observe, some topics sustain a relatively stable audience throughout the analysis period, while others exhibit a *sparking* behavior—characterized by sharp, short-lived increases in attention. To identify these spikes systematically, we compute a z-score for each topic’s daily distribution of chat participation, treating values exceeding three standard deviations above the mean as anomalous. We further filter out minor fluctuations by retaining only spikes involving more than 50 chats. This approach allows us to isolate significant bursts of attention likely triggered by 10 real-world events (see Table 1). Notably, during these *spark* periods, we observe a statistically significant

higher (Mann–Whitney U test) number of active chats (i.e., those exceeding ten daily messages), indicating that such episodic topics tend to ignite more intense and widespread engagement, not only within individual chats but also across the platform.

4.3 Thematic spheres and online niche communities

In the final stage of our analysis we identify *thematic spheres*, that is, groups of chats that consistently engage with the same topic on the same days. To do so, we first build a co-discussion network by connecting chats that have the same topic on the same day, and then apply a filtering approach to keep only edges reflecting high chance of coordination. This structure captures latent synchronized patterns of attention.

Once again, we extract the Giant Connected Component of this graph, which comprises 4,059 chats, and apply spectral clustering. The resulting partition reveals nine communities, with a modularity score of $Q = 0.490$, which suggests a well-defined community structure. Each community corresponds to a thematic sphere (Fig. 4 offers a qualitative network view). As shown in Fig. 5a, the size of these spheres is highly heterogeneous: while some communities account for a substantial fraction of the chats, with the largest sphere (community 4) containing 17.6% of all nodes, others are much smaller (with a minimum of 2.8% in community 8).

Figure 5b further characterizes each sphere by displaying the relative frequency of the most discussed topics within each community. Community 1 is strongly dominated by discussions revolving around conspiracy theories. These narratives also appear prominently in Community 5, although here they are combined with political breaking news, indicating a more hybrid focus. Community 3 is instead primarily engaged in discussions related to international conflicts and warfare. Communities 6, 7, and 8 are smaller in size and revolve around topics of self-improvement, the symbolic




Fig. 4 Qualitative visualization of the *Thematic Sphere* Graph, constructed from statistically significant instances of synchronized topic attention between chats. The green (n.2) and red (n.3) communities appear clearly separated, whereas the blue (n.0), orange (n.1), and brown (n.5) communities are more tightly connected, showing greater overlap in the topics they address

meaning of numbers and codes, and holistic or natural remedies.¹⁴ Notably, even in these peripheral spheres, conspiracy-related topics maintain a consistent presence, reflecting their cross-cutting nature. Community 4 shows a substantial presence of chats discussing antisemitic themes, often intertwined with broader political narratives. This suggests the existence of ideologically extreme and potentially neo-Nazi groups, whose discourse merges hate speech with conspiratorial and ultra-nationalist positions tied to contemporary political events. Community 2 is highly focused on cryptocurrency-related content, highlighting how financial narratives—despite being off-topic in a strict political sense—still circulate within the broader discussion ecosystem of political oriented chats. Finally, Community 0 emerges as the most politically cohesive sphere, encompassing discussions across a wide range of spark-triggered topics such as Trump’s guilty verdict, Biden’s withdrawal, and other institutional news. It stands out as the most tightly aligned with the core political agenda tracked in this study.

As defined in Section 3.2.4, we also build two complementary networks based on shared forwarded messages and common URLs to assess whether chats topically aligned also exhibit cohesive patterns of internal and external content circulation. Among the 4,059 chats involved in the topic-based network, 2,967 and 3,638 belong to the giant components of the internal and external content-sharing networks,

¹⁴ Despite being already visible at the message level, these topics were labeled as “in-target” based on a broader definition of political relevance, to avoid excluding content that might co-occur with political narratives in some chats. This choice proves meaningful, as these spheres—though not dominant—exhibit intermittent political engagement, especially during major events.

respectively. We refer to internal content as forwarded messages circulating within Telegram, and to external content as URLs linking to resources outside the platform. In both cases, we observe clear community structure, identifying 5 and 4 communities respectively (modularity scores $Q = 0.460$ and 0.464).

A clear structural alignment emerges from this analysis, as shown in Fig. 6. For instance, thematic sphere 4 shows the highest overlap (approximately 80%) with community 1 identified via forwarded messages, reinforcing the interpretation of this sphere as an ideologically cohesive and insular niche—characterized by consistent internal content circulation. Yet, when considering external content, sphere 4 overlaps most with sphere 3, suggesting that despite internal cohesion, it tends to draw from a broader pool of external sources shared with other communities. Similarly, while sphere 2 is thematically distinct, centered on cryptocurrency discussions, it shares forwarded messages with nodes from other spheres and shows its strongest URL-based alignment with sphere 0. This suggests once again that, although crypto-related discussions are off-topic, they remain intrinsically intertwined with the broader political debate.

Other spheres reveal more diffuse patterns. Sphere 5 shows weak overlap with any content-sharing community, consistent with a lower degree of structural cohesion and greater exposure to diverse information flows. Spheres 6 and 7, despite being topically distinct, show strong overlap in both forwarded and external content between each other, suggesting shared routines and a degree of isolation from the rest of the network.

More generally, URL-based communities appear less aligned with the structural isolation observed among thematic spheres. While forwarded content reinforces intra-sphere cohesion, URLs often circulate across spheres—e.g., between spheres 6 and 7, 1 and 8—suggesting a more fluid exchange of external resources. This may suggest that, while internal forwarding behaviors reinforce intra-sphere cohesion and thematic or ideological boundaries, external content (e.g., news sources, websites) crosses thematic boundaries more frequently and creates ties not captured by forwarding alone.

Overall, the interplay between synchronized topic attention and content-sharing behaviors reveals the multifaceted structure of Telegram communities. While topic alignment uncovers temporal convergence in discussion, content-sharing patterns expose deeper layers of community cohesion, shaped by both endogenous and exogenous information flows. Internal content, in the form of forwarded messages, tends to be highly specific to individual thematic spheres, reinforcing their distinct identity and cohesion. In contrast, external content such as URLs is more widely shared across spheres, often bridging otherwise disconnected communities. Taken together, our results provide evidence for the



(a) Community sizes.

(b) Most frequent topics.

Fig. 5 Overview of the community structure in the co-discussion network

existence of online niche communities within Telegram: clusters of chats that not only synchronize their attention around specific topics but also reinforce their cohesion through repeated and selective content circulation. However, this niche character is not uniform-some spheres exhibit strong insularity and alignment, while others engage in more porous and fluid exchanges across thematic and structural boundaries.

5 Discussion and conclusions

Our study presents a general-purpose, modular pipeline designed to support the analysis of conversational data on messaging platforms and social media. It is especially

suitable for high-volume data scenarios, combining semantic and structural signals to extract relevant content from corpora dominated by off-topic or spam-like material. While recent tools such as Telegram Monitor (Júnior et al. 2022), and the Telegram Observatory (Cavalini et al. 2023), have advanced monitoring capabilities for specific political contexts, they do not integrate structural and semantic modeling at daily resolution. By contrast, our pipeline enables both fine-grained temporal analysis and flexible application across platforms and thematic settings.

In this work, we apply the pipeline to analyze the Telegram ecosystem during the 2024 U.S. elections, with the goal of mapping the structure and thematic composition of political discussions across thousands of chats. While the design is tailored to the specific challenges of Telegram-such as weak moderation, multilingualism, and topical heterogeneity-its modular structure allows for adaptation to other platforms. The notion of a *chat* in our framework can be interpreted more broadly to encompass other conversational units, such as comment threads on Instagram posts, reply chains on X/Twitter, or discussions in YouTube channels. However, applying the method across platforms requires careful consideration of each platform’s interaction model, content longevity, and engagement patterns.

Crucially, the pipeline leverages the fine-grained semantic knowledge extracted from textual content through topic modeling, while also overcoming its limitations by integrating structural analysis via graph-based methods. Co-forwarding patterns allow us to infer the

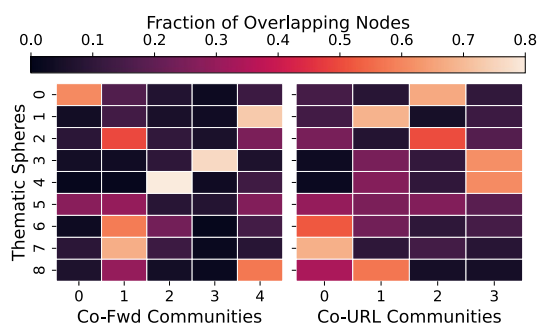


Fig. 6 Fraction of overlapping nodes among thematic spheres and the communities in the co-sharing of forwarded message and urls

thematic focus of low-volume or non-textual chats, while daily summarization and encoding produce coherent representations of ongoing discussions. While prior work favored either structural/network approaches (Venâncio et al. 2024; Cinus et al. 2025) or semantic/textual models (Paoletti et al. 2025; Vasconcelos et al. 2025; Gerard et al. 2025), our work bridges both approaches allowing for joint analysis of evolving topical attention and the underlying network structures that reinforce it.

The application of our pipeline to the `usc-tg-24-us-election` dataset provides an empirical answer to our first research question: prominent political topics during the 2024 U.S. elections emerge clearly in response to external triggers. We identify 36 refined chat-level topics, some driven by sustained engagement (e.g., war, conspiracy theories, etc.), others by event-driven spikes (e.g., Trump's assassination attempt or Biden's withdrawal). These bursts of attention are temporally localized and coincide with a tangible increase in overall platform activity. While previous research has documented topic polarization and curiosity-driven engagement (Vasconcelos et al. 2025), our contribution lies in capturing how specific communities on Telegram respond collectively to breaking events in real time, extending approaches like those in (Gerard et al. 2025) to broader, non-conflict-focused political discourse.

Our second research question addresses whether Telegram chats organize into thematic communities. The analysis of topic alignment reveals 9 distinct *thematic spheres*-sets of chats that repeatedly co-engage with the same topics on the same days. These spheres are topically coherent but also ideologically diverse: some revolve around antisemitic or far-right narratives, others around conspiracy theories, self-improvement, or alternative medicine. One sphere stands out for its moderate tone and alignment with institutional political discourse, while another focuses primarily on cryptocurrency-highlighting the entanglement between financial and political narratives on Telegram. This approach complements ideological mapping efforts like those by Alvisi et al. (2025), who examined large-scale ideological segmentation and toxicity in Italy. Our thematic spheres provide a dynamic lens into how such segmentation plays out across time, topics, and structural links within the chat ecosystem.

To investigate whether these thematic spheres also exhibit structural cohesion, we compare them with clusters derived from content-sharing patterns. Our results confirm that several spheres function as true *online niche communities*. The chats within these spheres do not merely discuss the same topics—they also share the same forwarded messages and, to a lesser extent, external URLs. Internal content (i.e., forwarded messages) proves highly specific to each sphere, reinforcing ideological or thematic boundaries. URLs, in turn, circulate more broadly across spheres, suggesting a looser structure of cross-topic

connections mediated by external sources. This layered content behavior echoes findings from Slobozhan et al. (2023), who observed differentiated roles of channels and groups during political protests in Belarus. Yet, our results show how this differentiation also manifests at the community level, with varying degrees of narrative cohesion and permeability.

This dual behavior illustrates the layered structure of Telegram's discourse: while thematic spheres are semantically cohesive, their content-sharing patterns vary in insularity. Some spheres display strong internal cohesion and isolation, typical of niche communities; others act as semi-permeable membranes, where thematic and structural boundaries are less sharply defined. By modeling both semantic alignment and content circulation, our work contributes to bridging the gap between message-level analyses (Hoseini et al. 2023; Kloos et al. 2024) and network-based models of political discourse (Venâncio et al. 2024; Gerard et al. 2025), offering a more integrated view of narrative evolution on Telegram.


5.1 Limitations and future works

Despite the robustness and scalability of our approach, several limitations must be acknowledged. First, our dataset truncates at the end of July 2024, constraining our ability to observe longer-term narrative developments during and after elections. Moreover, our filtering strategy focuses on English-dominant and politically salient chats, sacrificing breadth for depth and potentially omitting fringe or multilingual communities with relevant political content. The reliance on LLM-based summarization also introduces systematic artifacts when input quality is poor, occasionally generating spurious topics. Although mitigated through manual validation and topic refinement, this remains a structural limitation of large-scale LLM pipelines. Additionally, while instruction-tuned LLMs are effective at handling mainstream and especially Anglo-centric slang, they may fail to interpret highly localized or subcultural expressions, maybe unique to specific Telegram communities—an issue that affects even human annotators unfamiliar with the group context (Wuraola et al. 2024; Pavlick and Kwiatkowski 2019). To mitigate this, we guide the model to abstract away from stylistic markers and focus on salient discussion themes, accepting some loss in fidelity in exchange for semantic clarity. Future

Table 2 Dataset composition by entity type, language, and message volume

Entity Type	Count	English-dominant	Messages (M)	% of Messages
Groups	5,488	3,934	389.8	80.2%
Channels	23,398	13,726	96.5	19.8%
No messages	1,009	–	0.0	0.0%
Total	29,895	17,660	486.4	100.0%


Fig. 7 Distributions of activity and participation in Telegram chats



(a) Messages per chat.

(b) Users per group.

Fig. 8 Qualitative view of the Co-Sourcing Graph, sampling the 10% of the edges. Position of the nodes is computed through ForceAtlas2 (Jacomy et al. 2014)



(a) Co-Sourcing Communities.

(b) Dominant Languages (blue: English, green: German, red: Unknown).

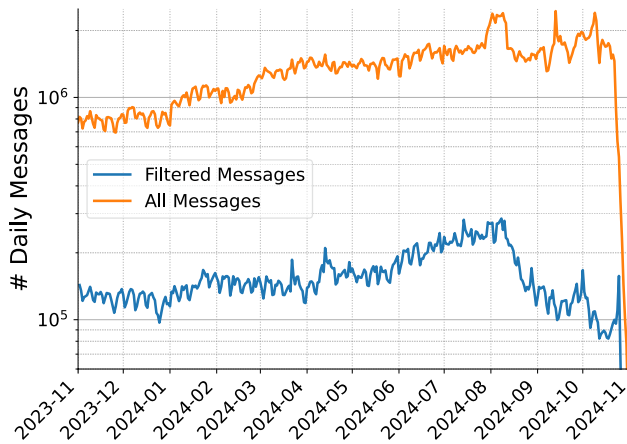


Fig. 9 Volume of messages over time. The orange line refers to all messages in the dataset; the blue one refers to the subset retained after filtering

work should address these challenges by incorporating multilingual topic detection and metadata-aware filtering, extending the time horizon of analysis, refining topic merging via more principled methods, and exploring strategies to better capture context-specific language. In addition, more flexible temporal segmentation—such as sliding or activity-adaptive windows—could help capture event-specific conversational dynamics beyond fixed daily intervals. Moreover, while our structural analysis reveals topical and content-based cohesion, it does not capture user-level coordination or influence. Extending this framework to include user dynamics, bot detection, and cross-platform propagation would offer a more comprehensive view of political discourse and manipulation strategies in messaging ecosystems.

Another interesting direction for future research concerns the role of external content—such as URLs—in shaping cross-sphere interactions. The fluidity observed in URL-sharing patterns suggests that, despite thematic




Fig. 10 Report of the LLM-based summarization process. The plots show the distribution of input token lengths, output lengths, reduction in tokens, processing time per summary, and number of message chunks per chat-day

segmentation, information may circulate across otherwise disconnected communities. Exploring how such exogenous flows influence opinion formation, exposure to alternative narratives, or community permeability would enrich our understanding of online debate dynamics.

Appendix

Dataset details

The distribution of message volume in the `usc-tg-24-us-election` dataset (Blas et al. 2025) is highly skewed, with a small number of chats—mostly groups—generating the majority of content; indeed, the top 10% most active chats generate 85.7% of the total message volume. A similar pattern holds for participation: the number of distinct users who actively send at least one message per group follows a heavy-tailed distribution. Table 2 summarizes the dataset composition. Figure 7a shows the distribution of message volume across all chats, highlighting the subset of messages sent specifically within groups. In contrast, Fig. 7b displays the number of active participants in the groups included in the dataset.

Supplementary material

Figure 8 shows a qualitative view of the co-sourcing graph displayed using a random 10% of the edges. Nodes represent chats and are colored to highlight their community (Fig. 8a) and their dominant language, i.e. the most used language in their messages (Fig. 8b).

Figure 9 shows the volume of messages over time, comparing the total volume with that remaining after careful filtering of the dataset (see Sec. 3.2.1).

Figure 10 shows key characteristics of the summarization task, including the distribution of input and output token lengths, the number of message chunks per chat-day, and the overall processing time. These statistics provide an operational overview of the scale and complexity of LLM-based summarization in our pipeline.

Acknowledgements This work has been partially supported by the Spoke 1 "FutureHPC & BigData" of ICSC — Centro Nazionale di Ricerca in High-Performance Computing, Big Data and Quantum Computing, funded by the European Union — NextGenerationEU. We also acknowledge the partial support from CNPq, CAPES, INCT-TILD-IAR, FAPEMIG, AWS, Google, NVIDIA, CIIASaúde, and FAPESP.

Author Contributions G.P. conducted the experiments. All authors contributed to the experimental design, interpreted the results, and wrote the manuscript.

Funding Open access funding provided by Politecnico di Torino within the CRUI-CARE Agreement.

Data Availability No datasets were generated or analysed during the current study.

Conflict of interest The authors declare no Conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alvisi L, Tardelli S, Tesconi M (2025) Mapping the Italian Telegram ecosystem: Communities, toxicity, and hate speech. arXiv e-prints pp arXiv:2504
- Anglano C, Canonico M, Guazzone M (2017) Forensic analysis of telegram messenger on android smartphones. Digit Investig 23:31–49
- Barnes K, Juhász P, Nagy M, et al (2024) Topicality boosts popularity: a comparative analysis of NYT articles and Reddit memes. SNAM 14(1)
- BehnamGhader P, Adlakha V, Mosbach M, et al (2024) Llm2vec: Large language models are secretly powerful text encoders. arXiv:2404.05961
- Blas L, Luceri L, Ferrara E (2025) Unearthing a billion Telegram posts about the 2024 US presidential election: Development of a public dataset. In: Proc. Comp. of The Web Conf '25
- Buzelin A, R. Dutenhefner P, S. Locatelli M, et al (2025) Analyzing political discourse on Discord during the 2024 US presidential election. In: Proc. ACM WebSci '25
- Cavalini A, Malini F, Gouveia F et al (2023) Politics and disinformation: analyzing the use of telegram's information disorder network in Brazil for political mobilization. First Monday. <https://doi.org/10.5210/fm.v28i5.12901>
- Chagas V (2022) WhatsApp and digital astroturfing: a social network analysis of Brazilian political discussion groups of Bolsonaro's supporters. Int J Commun 16
- Cinus F, Minici M, Luceri L, et al (2025) Exposing cross-platform coordinated inauthentic activity in the run-up to the 2024 us election. In: Proc.The Web Conf'25
- Dall'Amico L, Couillet R, Tremblay N (2021) Nishimori meets Bethe: a spectral method for node classification in sparse weighted graphs. Theory Exp, J Stat Mech
- Dieng AB, Ruiz FJ, Blei DM (2020) Topic modeling in embedding spaces. Transactions of the Association for Computational Linguistics 8
- Dong Z, Tang T, Li L, et al (2023) A survey on long text modeling with transformers. arXiv:2302.14502
- Ferreira CHG, Murai F, Silva APC, et al (2021) On the dynamics of political discussions on Instagram: A network perspective. OSNAM 25
- Gerard P, Volkova S, Penafiel L, et al (2025) Modeling information narrative evolution on Telegram during the Russia-Ukraine war. In: Proc. ICWSM'25, pp 602–614

- Grattafiori A, Dubey A, Jauhri A, et al (2024) The LLama 3 Herd of Models. [arXiv:2407.21783](https://arxiv.org/abs/2407.21783)
- Grootendorst M (2022) BERTopic: Neural topic modeling with a class-based TF-IDF procedure. [arXiv:2203.05794](https://arxiv.org/abs/2203.05794)
- Guo Y, Wang D, Wang L, et al (2024) Beyond app markets: Demystifying underground mobile app distribution via Telegram. *Proc ACM Meas Anal Comput Syst* 8(3)
- Hanley HW, Durumeric Z (2024) Partial mobilization: Tracking multilingual information flows amongst Russian media outlets and Telegram. In: *Proc. ICWSM'24*
- Herasimenka A, Bright J, Knuutila A et al (2023) Misinformation and professional news on largely unmoderated platforms: The case of Telegram. *J Inf Technol Polit* 20(2):198–212
- Hoseini M, Melo P, Benevenuto F, et al (2023) On the globalization of the QAnon conspiracy theory through Telegram. In: *Proc. ACM WebSci'23*
- Hoyle A, Goel P, Peskov D, et al (2021) Is automated topic model evaluation broken? the incoherence of coherence. In: *Proc. NeurIPS'21*
- Jacomy M, Venturini T, Heymann S et al (2014) ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PLoS ONE* 9(6):1–12. <https://doi.org/10.1371/journal.pone.0098679>
- Júnior M, Melo P, Kansao D, et al (2022) Telegram monitor: monitoring Brazilian political groups and channels on Telegram. In: *Proc. ACM Hypertext'22*
- Khandelwal T (2024) Investigating the impact of text summarization on topic modeling. [arXiv:2410.09063](https://arxiv.org/abs/2410.09063)
- Kloo I, Cruickshank IJ, Carley KM (2024) A cross-platform topic analysis of the nazi narrative on Twitter and Telegram during the 2022 Russian invasion of Ukraine. In: *Proc ICWSM'24*
- La Morgia M, Mei A, Mongardini AM (2025) TGDataset: collecting and exploring the largest Telegram channels dataset. In: *Proc. ACM SIGKDD'25*
- Liu Y, Ott M, Goyal N, et al (2019) Roberta: a robustly optimized BERT pretraining approach. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692)
- Maros A, Almeida JM, Benevenuto F, et al (2020) Analyzing the use of audio messages in WhatsApp groups. *Proc The Web Conf'20*
- Mimno D, Wallach H, Talley E, et al (2011) Optimizing semantic coherence in topic models. In: *Proc. EMNLP '11*
- Muennighoff N, Tazi N, Magne L, et al (2023) MTEB: Massive text embedding benchmark. In: *Proc. EACL '23*
- Nakatani S (2010) `langdetect`. <https://pypi.org/project/langdetect/>. Accessed 12 May 2025
- Naseri M, Zamani H (2019) Analyzing and predicting news popularity in an instant messaging service. In: *Proc. ACM SIGIR'19*
- Nikkhah S, Miller AD, Young AL (2018) Telegram as an immigration management tool. In: *Companion Proc. ACM CSCW '18*
- Nobre GP, Ferreira CH, Almeida JM (2022) A hierarchical network-oriented analysis of user participation in misinformation spread on WhatsApp. *Inf Process Manag*
- OpenAI (2024) GPT-4o mini: advancing cost-efficient intelligence. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>. Accessed 29 May 2025
- Paoletti G, Dall'Amico L, Kalimeri K, et al (2024) Political context of the European vaccine debate on Twitter. *SciRep* 14(1)
- Paoletti G, Almeida JM, Vassio L, et al (2025) Join the chat: How curiosity sparks participation in Telegram groups. In: *Proc. ICWSM'25*
- Pavlick E, Kwiatkowski T (2019) Inherent disagreements in human textual inferences. *Trans Assoc Comput Linguist* 7:677–694
- Perlo A, Paoletti G, Jha N, et al (2025) Topic-wise exploration of the Telegram group-verse. In: *Proc. Comp. ACM The Web Conf.'25*
- Resende G, CS Reis J, Vasconcelos M, et al (2019) Analyzing textual (mis) information shared in WhatsApp groups. In: *Proc ACM WebSci'19*
- Rousseeuw PJ (1987) Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J Comput Appl Math* 20
- Slobozhan I, Brik T, Sharma R (2023) Differentiable characteristics of Telegram mediums during protests in Belarus 2020. *SNAM* 13(1):19
- Thorgren E, Mohammadinodooshan A, Carlsson N (2024) Temporal dynamics of user engagement on Instagram. In: *Proc. ACM WebSci '24*
- Urman A, Katz S (2022) What they do in the shadows: examining the far-right networks on Telegram. *Inf Comm Soc* 25(7):904–923
- Vasconcelos FF, Ferreira CH, Magno de Sousa A, et al (2025) The role of curiosity in information dissemination on Telegram: a case study from the 2022 Brazilian elections. In: *Proc. ACM WebSci '25*
- Venna J, Kaski S (2005) Local multidimensional scaling with controlled tradeoff between trustworthiness and continuity. In: *Proc. WSOM'05*
- Venâncio OR, Ferreira CH, Almeida JM, et al (2024) Unraveling user coordination on Telegram: a comprehensive analysis of political mobilization during the 2022 Brazilian presidential election. In: *Proc. ICWSM'24*
- Walther S, McCoy A (2021) US extremism on telegram. *Perspect Terror* 15(2):100–124
- Wich M, Gorniak A, Eder T, et al (2022) Introducing an abusive language classification framework for Telegram to investigate the German hater community. In: *Proc. ICWSM'22*
- Wuraola I, Dethlefs N, Marciniak D (2024) Understanding slang with LLMs: Modelling cross-cultural nuances through paraphrasing. In: *Proc. EMNLP '24*, pp 15525–15531
- Yang Y, Paudel R, McShan J et al (2025) Coordinated link sharing on facebook. *SciRep* 15(1):15684

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.