Analyzing Polarization in Online Discourse on the 2023-2024 Israel–Hamas War

Daniel Miehling, Daniel Dakota, Sandra Kübler
Indiana University
{damieh, ddakota, skuebler}@iu.edu

Abstract

We investigate large-scale sentiment analysis of YouTube Shorts in the context of the 2023-2024 Israel–Hamas war. Using a corpus of over 3 million user comments and replies from four state-funded or state-supported international media channels, we track sentiment toward key geopolitical and ideological entities over the course of one year. We investigate the correspondence between large-scale sentiment analysis and manual fine-grained analysis of trends as well as the correspondence between a longitudinal analysis and geopolitical events. Results show that overall sentiment trends depict user attitudes, but to interpret these patterns correctly, we need domain-specific knowledge and a combination of the automatic analysis with fine-grained manual analysis. We also show that the peaks of a longitudinal analysis correspond to (geo-)political events such as the Eurovision Song Contest.

1 Introduction

The October 7, 2023 attacks and the ensuing war between Israel and Hamas triggered a surge in global online discourse characterized by intense polarization and ideological fragmentation. While polarization in digital spaces has been widely studied, most analyses of political communication rely on coarse methods such as stance detection or user clustering, often applied to Twitter (Becatti et al., 2019; Stier et al., 2018). In contrast, we focus on using large-scale aspect-based sentiment analysis to investigate trends in user-generated content.

More specifically, we use aspect-based sentiment analysis (ABSA) to analyze over 3 million comments and replies on YouTube Shorts (a short-form video format similar to TikTok videos) by four state-funded or state-supported international broadcasters (Cull, 2008) (Al Jazeera, TRT World, BBC News, and Deutsche Welle)[1]. Since ABSA enables the detection of distinct sentiment toward multiple entities within the same utterance (Chauhan and Meena, 2019; Mai and Le, 2021), it can offer a nuanced view of ideological alignment. Consider the following example[2]:

(1) *Hamas being labeled as terrorist group is just a label. They are resistances against occupation and atrocities... Some countries like Israel and America have caused more dead a terror... but you can't label them Terrorists because they are countries.*

Here, *Hamas* is framed in a relativizing tone that challenges its designation as a terrorist group. In contrast, *Israel* and *America* are depicted as more violent actors, invoking negative sentiment. The final mention of *Terrorists* emphasizes a perceived double standard, adding a layer of ideological critique. Thus, ABSA allows these overlapping views to be disentangled.

Our study applies a current ABSA model (Yang et al., 2021, 2023) to investigate political communication in the digital mainstream. Drawing on insights from Kušen and Strembeck (2023), who show how emotional exposure during the early stages of the Ukraine war shaped user affect over time, we use large-scale sentiment analysis to map fine-grained sentiment toward competing and often opposing political actors. This approach enables us to track how emotional alignment and ideological positioning shift over time within digital discourse on YouTube. We trace sentiment trajectories over a 12-month period following the Hamas-led attack on Israel. We aim to examine the correspondence of sentiment analysis with fine-grained manual analysis as well as the correspondence to spikes in the longitudinal analysis to (geo-)political events such

[1] In our tables and figures, we use the abbreviations BBC and TRT to refer to BBC News and TRT World, respectively.

[2] The example represents user-generated content drawn from our corpus.

Proceedings of the 21st Conference on Natural Language Processing (KONVENS), Volume 2: Workshops, pages 7–16, 2025

as the Eurovision Song Contest. More specifically, we ask the following research questions:

- **RQ1:** To what extent does large-scale sentiment analysis depict trends in user-generated content on Shorts videos?
- **RQ2:** How well do sentiment trends in a longitudinal analysis correspond to geopolitical events?

2 Related Work

With extensive prior research examining hate speech, misinformation, and political extremism in online spaces (Becker et al., 2023; Brown et al., 2024; Finkelstein et al., 2023; Rieger et al., 2020; Topor, 2024), sentiment analysis has emerged as a tool to model ideological alignment, polarization, and user engagement in socio-political discourse.

To capture the implicit and coded nature of political speech, Subramanian et al. (2023) and Young et al. (2024) emphasize the importance of context-sensitive sentiment models capable of interpreting nuance in polarized discourse. Shifts in media sentiment can precede conflict outbreaks (Jamison et al., 2023), with sentiment dynamics reflecting affective polarization (Lerman et al., 2024), and can capture emotional reactions in crisis moments (Kušen and Strembeck, 2023; Win Myint et al., 2024).

Aspect-based sentiment analysis has proven suitable for modeling longitudinal sentiment variation in politically charged online discussions. It captures more nuanced sentiment toward discrete entities or themes, enabling robust sentiment extraction at the aspect level across multiple targets within the same comment (Rietzler et al., 2020; Zhang et al., 2022), particularly in domains such as hate speech detection (Mughal et al., 2024; Zainuddin et al., 2016, 2018; Zhang et al., 2024) and political communication (Gold et al., 2018; Miok et al., 2023; Seno et al., 2024).

3 Methodology

3.1 Corpus and Collection

The empirical data was obtained via the official YouTube API[3]. We started the corpus collection two weeks after the events of October 7. To narrow the scope of this study, we focus on state-funded or state-supported international outlets: *TRT World*

[3] https://developers.google.com/youtube/v3/docs

Outlet	Videos	Comments	Users
DW	84	9 172	6 521
AJ	958	1 187 470	388 384
BBC	70	23 039	13 517
TRT	1 258	2 224 981	773 773
Total	2 370	3 444 662	787 157

Table 1: Corpus statistics across media outlets.

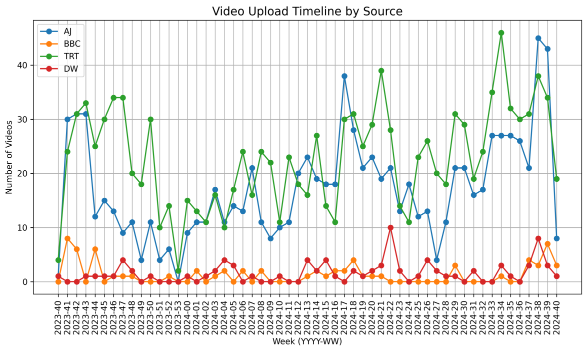

Figure 1: Weekly frequency of YouTube Shorts uploads per outlet over a 12-month period.

(Turkey), *BBC News* (United Kingdom), *Al Jazeera* (Qatar), and *Deutsche Welle* (Germany)[4]. These channels play a critical role in shaping public discourse and often reflect distinct geopolitical perspectives. Over a twelve-month data collection period, we gathered more than 3.4 million comments and replies from these four outlets, generated by over 780 000 unique users. Table 1 shows the total number of unique video IDs, user-generated comments and replies, and distinct users participating in discussions for each outlet.

Figure 1 illustrates the weekly frequency of YouTube Shorts uploads per source. *TRT World* and *Al Jazeera* consistently released substantially more content, often exceeding 30 to 40 Shorts per week, compared to *BBC News* and *DW*, which uploaded fewer than 10 Shorts weekly. *BBC News* and *DW*, in turn, cover a broader thematic spectrum, with fewer videos directly addressing the Israel-Gaza conflict. While we initially considered including earlier data, a subsequent rescraping attempt on user-generated content related to BBC's Shorts revealed that nearly 50% of the original comments and replies had been removed, likely due to YouTube's content moderation policies.

[4] While *TRT World* has stated that it is not affiliated with the Turkish government, the other three channels are state-funded or state-supported broadcasters.

Outlet	Overall	Non-Eng.	Non-Eng. %
TRT	771 855	149 268	19.33
AJ	280 744	47 032	16.75
DW	2 733	386	14.12
BBC	10 298	850	8.25

Table 2: Remaining content after language filtering.

3.2 Preprocessing

We implemented a multi-step preprocessing pipeline to prepare the corpus for sentiment analysis. First, we filtered structurally invalid JSON files and removed entries with missing fields. We then normalized user mentions (e.g., `@username`) to support accurate tokenization. To ensure language consistency, we applied the `langdetect` library[5] and retained only English-language comments. For further quality control, we computed an English lexical coverage ratio using the `nltk.corpus.words` vocabulary[6] and discarded comments with less than 40% English terms (see Table 2).

To constrain the analysis to relevant discourse, we filtered the corpus using a targeted aspect lexicon that included geopolitical and ideological entities central to the conflict. This lexicon comprised the following terms: *Israel*, *Palestine*, *Jews*, *Palestinians*, *Zionists*, *Hamas*, *Hezbollah*, and *Muslims*[7].

3.3 Aspect-Based Annotation

Manual annotations were necessary to fine-tune an aspect-based sentiment analysis model to annotate the large scale corpus. The user-generated content in our corpus contains considerable noise, including unconventional spellings, typographical errors, informal language, and domain-specific allusions or "dog whistles". To ensure accurate sentiment classification, we used a small subset of our corpus and annotated it for sentiment. We used `Label Studio`[8] (Tkachenko et al., 2025), an open-source data labeling platform, to annotate text segments with both aspect categories and sentiment labels. The annotation was conducted by an expert in the field. Sentences were pre-selected using dependency parsing (see Section 3.4) to identify those containing aspect terms from the aforementioned lexicon. Sentiment was annotated as *positive*, *neutral*, or *negative*, and applied to both explicit and clearly implied references. The annotation scheme accounted for informal syntax, metaphorical language, and ideological cues, which are characteristic of user-generated discourse.

Aspect	Negative	Neutral	Positive
Hamas	576	422	225
Hezbollah	5	14	12
Israel	296	278	376
Jews	192	244	109
Muslims	112	272	234
Palestine	175	186	101
Palestinians	272	261	260
Zionists	228	272	25
Total	1 856	1 949	1 342

Table 3: Distribution of annotated segments by aspect and sentiment label.

The final annotated corpus comprises 5 147 text segments[9], sampled from comments and replies across the four state-funded or state-supported media outlets. Selection criteria included lexical variety, comment diversity, and the removal of near-duplicate or semantically redundant content. We aimed to balance the sentiment distribution within each aspect category, although certain entities, particularly *Zionists*, were predominantly associated with negative sentiment, limiting class balance. Table 3 presents the distribution of annotated segments by aspect and sentiment label.

To assess reliability, a stratified random sample of 500 segments was independently annotated by a second expert. We computed inter-annotator agreement (IAA) on this subset, yielding scores of $\kappa = 0.86$ and $\alpha = 0.85$, indicating substantial agreement (see Table 4). While these scores reflect consistency, they do not guarantee conceptual accuracy in highly polarized or coded language domains (Paun et al., 2022), where ideological framings and implicit sentiment pose inherent challenges.

3.4 Dependency Parsing

We used the biaffine graph-based dependency parser (Dozat and Manning, 2017) implemented in

[5]https://github.com/Mimino666/langdetect

[6]https://www.nltk.org/_modules/nltk/corpus/reader/wordlist.html

[7]After processing, the aspect *Hezbollah* appeared only in a small number of examples and was therefore excluded from the ABSA evluation due to insufficient data.

[8]https://labelstud.io

[9]Each segment corresponds to a dependency-parsed sentence or clause containing at least one aspect mention. Due to informal punctuation, ellipses, or run-on constructions, some parsed segments span multiple clauses.

Sentiment	Cohen's κ	Krippendorff's α
Negative	0.850	0.850
Neutral	0.867	0.868
Positive	0.869	0.869
Overall	0.862	0.855

Table 4: Inter-annotator agreement scores across sentiment classes.

Class	Prec.	Rec.	F1
Negative	73.5	84.4	78.6
Neutral	78.4	71.9	75.0
Positive	85.0	77.9	81.3
Avg.	78.4	78.0	77.9

Table 5: Overall sentiment classification performance (on the test set).

the SuPar library[10] to extract syntactic structures, identify aspect terms, and locate their grammatical heads. In our workflow, the parser detected dependency relations within each comment, allowing us to align syntactic heads with predefined aspect categories for aspect term extraction (Zhang et al., 2022). Parsing was conducted using the pretrained `biaffine-dep-en` model, trained on English Universal Dependencies treebanks. We performed inference on a GPU with batch-wise processing over tokenized user-generated text.

3.5 Model Training and Evaluation

For aspect-based sentiment analysis, we used the end-to-end `DeBERTa-v3-large-absa-v1.1` model provided in the `pyABSA` library (Yang et al., 2021, 2023). This pretrained model was initially trained on English-language benchmark datasets, including Twitter and SemEval corpora, for aspect-based sentiment classification. We finetuned the model for 5 epochs on our task-specific training set using an NVIDIA A100 GPU. Training used a batch size of 2 and gradient accumulation over 8 steps, resulting in an effective batch size of 16. A cosine learning rate scheduler was used with an initial learning rate of 1×10^{-5}. As no separate development set was defined, we selected the best model checkpoint based on validation loss. Following the aspect-prompted classification framework, aspect terms in each sentence were replaced with a `$T$` marker. Input sequences were truncated to a maximum length of 512 tokens.

Table 5 shows the evaluation on the test set (precision, recall, and F1 across sentiment classes). The model achieves a macro-averaged F1 of 77.9, indicating strong performance in classifying explicit sentiment expressions related to our aspect terms.

The model most frequently confuses neutral comments and negative ones in both directions (see Figure 2). Since our goal is to assess sentiment

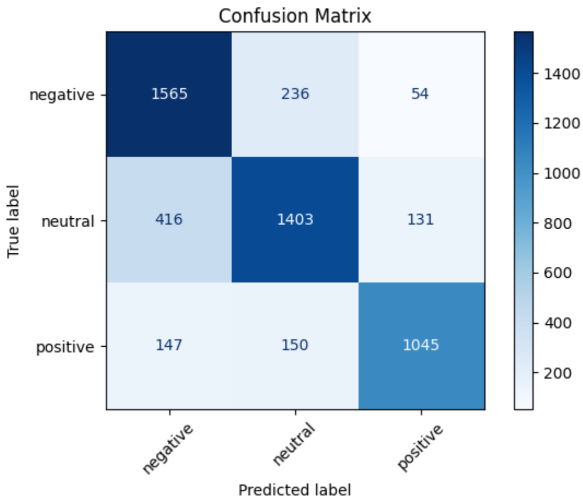

Figure 2: Confusion matrix for sentiment classification (on the test set).

directed toward specific ideological actors, we further evaluated per-aspect performance. As shown in Table 6, model performance is consistent across aspects, with particularly high accuracy for *Hezbollah*, *Muslims*, and *Palestine*, and solid results for more contentious aspect terms such as *Zionists* and *Hamas*. Our error analysis revealed three primary challenges: (1) difficulty in detecting implicit sentiment, particularly in sarcastic or coded language; (2) misclassifications in sentences containing ambiguous sentiment toward multiple aspects (e.g., "*Free Jews from Israel*"); and (3) challenges in processing code-switching comments and transliterations, where sentiment-bearing terms appeared in hybrid forms (e.g., Arabic-English transliterations).

4 Findings

To address our research questions, we adopt a two-step analytical approach that combines quantitative analysis with close qualitative reading. While ABSA enables large-scale sentiment classification, many findings cannot be interpreted in aggregation alone. In response to RQ1, we begin by analyz-

[10] https://github.com/yzhangcs/parser

Aspect	Acc.	Prec.	Rec.	F1
Hamas	75.1	75.0	75.1	74.6
Hezbollah	90.3	90.6	90.3	90.3
Israel	75.4	76.6	75.4	75.6
Jews	74.3	74.9	74.3	74.3
Muslims	85.4	85.8	85.4	85.5
Palestine	81.6	81.9	81.6	81.5
Palestinians	78.9	78.9	78.9	78.7
Zionists	79.0	79.6	79.0	78.7

Table 6: Per-aspect sentiment classification performance (on the test set).

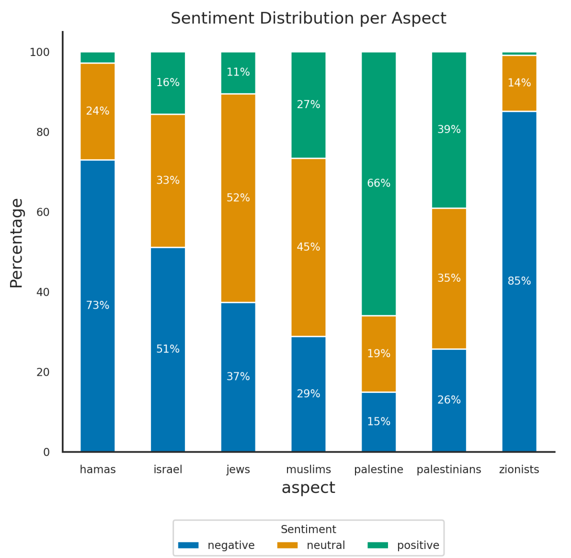

Figure 3: Stacked bar chart for sentiment distribution per aspect (in %).

ing the aggregated sentiment distribution across aspect terms to uncover general patterns of political communication in user-generated reactions to Shorts. We then turn to RQ2 by examining longitudinal sentiment dynamics and selected discourse patterns that reveal ideological alignments amid heightened volumes of user-generated content.

4.1 Aggregated Sentiment Distribution

Figure 3 presents the aggregated results of the sentiment analysis in form of the distribution of sentiment labels (negative, neutral, positive) across all classified aspect terms in the corpus, aggregated over all four media outlets. Negative sentiment dominates for the aspects *Israel*, *Zionists*, and, to a lesser extent, *Jews*. In contrast, *Palestinians* and *Muslims* receive more mixed or neutral sentiment, reflecting framings across the analyzed outlets. Interestingly, *Zionists* receive a higher proportion of negative sentiment (85%) than *Hamas* (73%), indicating that references to Zionism in mainstream YouTube discourse may elicit particularly adversarial responses[11].

While the aggregated results indicate a measurable share of positive sentiment toward *Palestine* and *Palestinians*, a closer examination of sentiment associated with *Jews* reveals that positive expressions do not necessarily reflect broadly favorable attitudes. Instead, such sentiment often appears in contrastive framings, for example, praising "good Jews" who oppose Israel or Zionism, while implicitly or explicitly condemning others, see examples (2) and (3). This discursive pattern reflects a form of conditional inclusion, where selective approval reinforces ideological boundaries. As mentioned earlier, such patterns of selective framing are not captured by sentiment proportions alone and require close reading to reveal how users shape ideological narratives.

(2) There are plenty of Jewish people fighting for Palestinian independence, we should not conflate Jewish people with the genocidal state of Israel.

(3) There are good orthodox Jews but Israel is run by a sadistic Zionist

A similar ambiguity emerges in references to *Palestinians*. In example (4), the model assigns positive sentiment, yet the phrase carries an implicit reference to a violent act (i.e., the "pager attack" linked to Hezbollah), thus revealing antagonistic intent beneath a lexically positive surface[12].

(4) Free Pagers for Palestine

Despite such conceptual ambiguities, these examples shed light on how evaluative language in user-generated content often carries multi-layered or contrastive meanings. In example (5), the phrase "Free Palestine" appears lexically positive, yet it is reframed as a critique of Hamas. Such contrastive uses illustrate how sentiment-laden expressions may simultaneously convey positive alignment with one political actor while denouncing

[11] While the proportion of negative sentiment is higher for *Zionists*, the total number of classified instances is much lower: 14 358 (*Zionists*) vs. 171 114 (*Hamas*).

[12] The phrase refers to events on 17–18 September 2024 in Lebanon, where explosives concealed in pager devices were deliberately detonated to target members of Hezbollah.

another. The presence of multiple aspect terms within short, syntactically simple comments suggests that such juxtapositions are common in political discourse and reveal how users employ sentiment cues to express ideological alignment in response to Shorts.

(5) Free Palestine from Hamas

While the majority of sentiment directed at *Hamas* is negative, we identified a smaller subset of positive sentiment. These instances often reflect implicit endorsement of Hamas as a resistance actor, as illustrated in example (6) and example (7).

(6) Victory to HAMAS-HEZBOLLAH FREE PALESTINE

(7) Kill every israel soldier who attack on falastin, hamas

Frequently, intensified sentiment is conveyed through symbolic and visual elements. Emoji sequences often accompany slogans like "Free Palestine" or denunciations of Zionism, that may serve to amplify affective tone or signal ideological stance. As shown in example (8), semiotic markers are combined to emphasize solidarity and resistance.

(8) FREE PALESTINE NOW! ✊🏾 PS✌🏾 Down with 💤ionism and apartheid 🤮 🤬

Across the corpus, emojis are used extensively: 18.1% of all classified sentences (181 095 out of 1 000 750) contain at least one emoji. This high frequency highlights their role as affective intensifiers in discourse and underscores the need to consider multimodal cues as crucial elements of contemporary political communication within the digital mainstream. However, interpreting these patterns often requires domain-specific knowledge, as their meanings can diverge significantly from conventional language use.

4.2 Identifying Longitudinal Dynamics

To analyze how shifts in user discourse are reflected in sentiment patterns, we perform a peak analysis of weekly aspect-based sentiment trends. Figure 4 shows the overall frequency of sentiment-labeled content across the four state-funded or state-supported international broadcasters. It reveals key differences in user engagement: for instance, Shorts on *TRT World* consistently resulted in high volumes of user-generated content, while activity

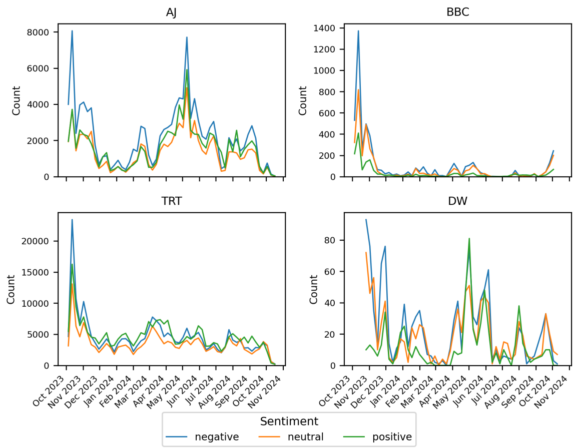

Figure 4: Sentiment trends over time across outlets.

on *DW* remained sporadic, only increasing in response to specific events. We observe repeated spikes in negative sentiment that coincide with geopolitical events such as the Eurovision Song Contest or Israel's Independence Day. However, instead of focusing solely on these general trends, we identify discursive tipping points: high aggregations of emotionally charged user responses that intensified user-generated reactions. We examine these sentiment surges more closely through aggregated data of weekly aspect-based sentiment trends, which allows us to trace how user sentiment fluctuated in response to unfolding real-world events.

Figure 5 shows sentiment trends per aspect. This analysis reveals three major peaks between March and May 2024, corresponding to heightened sentiment around *Palestine*, *Israel*, and *Zionists*. For each peak, we retrieved the most-commented Shorts videos and analyzed their content in relation to user responses. We compared sentiment curves across outlets and examined video titles and descriptions, together with comments and replies, to assess how user-generated activity reflects reporting on geopolitical events.

The first peak in early May coincides with the release of Shorts covering the Eurovision Song Contest, see example (9), and international campus protests, see example (10). High-engagement videos contributing to this spike included footage of Dutch riot police dismantling a Gaza protest camp and scenes of pro-Palestinian chants during Israel's Eurovision rehearsal.

(9) Well done, to the audience at the Eurovision. Letting Israel know how the world feels about them !! FREE PALESTINE NOW !!!

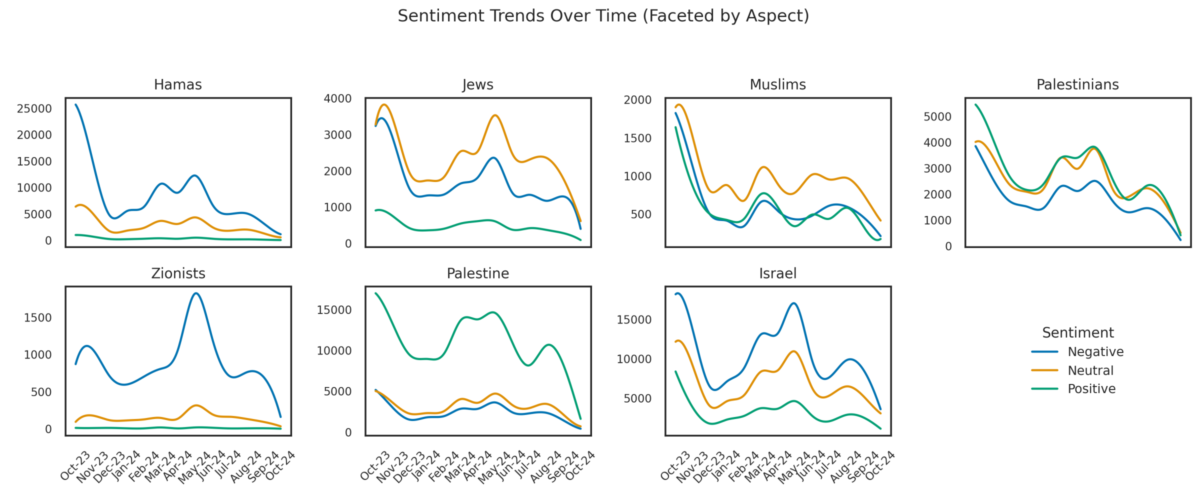

Figure 5: Sentiment trends over time across seven aspects.

(10) Many thanks the students who support majlum Palestine

These videos triggered waves of solidarity expressions such as “Free Palestine”, as well as accusations of genocide, see example (11), illustrating how users signaled political stance.

(11) Just need humanity to stand with Palestine. Stop genocide

The second peak coincided with Israel’s Independence Day on May 14 and was shaped by Shorts that cover allegations of the International Criminal Court (ICC) against Israel, as well as a Short depicting US-supplied weapons used in Gaza. During this period, negative sentiment toward *Israel* intensified. Users invoked narratives of indigeneity and dispossession, see examples (12) and (13), reflected by a high frequency of trigrams such as “Canaan Palestinians descendants” and “indigenous people Palestinians” to assert historical and moral claims.

(12) Wasn’t Israel made up by the whiteEuropeans, Delete the Hebrews field of land from the indigenous canaanites?

(13) The indigenous people are the Palestinians not the israeles who’ve been separated from Canaan and the old israel for a very long time

A third peak followed shortly thereafter, marked by intensified negative sentiment toward *Zionists*. This spike coincide with Shorts depicting police crackdowns on protest camps and a widely viewed video showing a scholar harassing a Muslim woman at a pro-Israel protest. Commenters responded with ideologically charged language, adopting anti-colonial framings and linking protest repression to broader narratives of systemic violence. Phrases such as “Zionists are dangerous”, see example (14), and “Zionist colonizers”, see example (15), highlight how the discourse shifted from solidarity-based expressions to overt ideological opposition.

(14) Zionists are dangerous for all people of the world.

(15) You mean Zionists hate them spreading the truth and standing up against the evil Zionist colonizers committing genocide.

Compared to earlier peaks, the third peak features significantly more pejorative and antagonistic language. Users adopt moral binaries and dehumanizing rhetoric to frame the conflict between Hamas and Israel, shifting from solidarity-based expressions to ideologically charged attacks targeting *Zionists*, marking a clear intensification of discursive polarization.

Taken together, the comparison across these peaks suggests a transition from expressions of solidarity to increasingly polarizing and antagonistic rhetoric. User-generated content toward *Zionists* and *Israel* exhibits more sentiment shifts than content related to *Palestine*. Users responded to geopolitical flashpoints by asserting legitimacy claims, voicing collective grievances, and expressing identity-based solidarities, dynamics consistent with prior research on event-driven polarization (Royesh and Grossman, 2021; Miehling, 2024; Alamsyah et al., 2024; Kušen and Strembeck, 2023).

5 Discussion

Our findings suggest that sentiment in user-generated content functions not merely as an emotional signal, but as a means of ideological expression and alignment. ABSA enables us to trace how users position themselves in relation to geopolitical actors through coded language and contrastive framings.

These patterns were especially salient in discourse surrounding *Zionists* and *Israel*, which were frequently associated with accusations of colonialism, genocide, and systemic violence. In contrast, while the aspect *Hamas* displayed only a small share of positive sentiment, its occasional framing in terms of resistance or heroism illustrates how affective discourse can operate through implication and conceptual ambiguity rather than overt endorsement. In particular, *Zionists* received the highest proportion of negative sentiment, even exceeding *Hamas*. While this trend underscores adversarial positioning toward Zionism as an ideology, the lack of positive examples also posed challenges for classification. To maintain empirical transparency, and to avoid biasing the ABSA system, we refrained from augmenting the training data, e.g., by generating synthetic data, to address class imbalance.

At the same time, our findings demonstrate the utility of large-scale sentiment analysis not only for identifying sentiment polarity, but also for unpacking rhetorical strategies users deploy to articulate political stances. In this context, affect operates discursively, conveying moral evaluation and ideological positioning. We observed persistent ambiguity in contrastive expressions involving Jews, where seemingly positive portrayals of Jewish fringe groups often legitimized the simultaneous condemnation of Israel or *Zionists*. Such constructions not only complicate the interpretation of straightforward sentiments but also reflect deeper ideological divisions.

Temporal sentiment patterns in our peak analysis further reveal that user-generated content extends beyond immediate reactions to breaking news, popcultural events, or protest repression. Affective intensification often mirrors geopolitical flashpoints but also transcends them. The accompanying discourse frequently moves beyond the specific Short-video context. Users employ emotionally charged, morally coded language to express solidarity, outrage, and condemnation. Thus, user-generated content should not be understood solely as real-time responses, but as discursive echoes, shaped by collective memory and entrenched resentment. Taken together, these findings suggest that sentiment analysis leads to meaningful results for interpreting political discourse online, not as a series of isolated opinions, but as structured, emotionally mediated ideological formations.

6 Conclusion and Future Work

We have investigated the interconnection of sentiment and polarization in political discourse. To answer RQ1, we have demonstrated that large-scale sentiment analysis can capture patterns of political communication embedded in user-generated discourse. While surface-level trends reveal general sentiment, a closer reading of aspect-specific results, such as contrastive framings (“good Jews” vs. “Zionists”), semiotic markers, and selective sentiment toward ideologically loaded terms, shows how users employ affective language to articulate political stance, moral judgment, and group alignment.

Addressing RQ2, we showed that we can identify evolving discourse patterns and ideological shifts in response to real-world geopolitical events. By analyzing sentiment peaks and comparing aspect-specific dynamics, we traced users expressing both solidarity and polarization, revealing affective intensification and discursive realignment over time.

For future work, we will extend our investigation and analyze more sentiment targets (aspects). We will also investigate the usability of an RST parser to gain deeper insights into how the discourse is structured beyond the sentence level.

Limitations

Parsing: While dependency parsing enhanced aspect-term extraction, it did not fully resolve ambiguities in sentiment attribution. Future improvements may benefit from incorporating contrastive learning techniques or domain-adapted embeddings to better capture context-sensitive sentiment in highly polarized political discourse.

Annotations: Some statements lack clear intent markers (e.g., sarcasm, irony, rhetorical questions). We infer meaning based on domain-specific expertise using established patterns of dog whistles and coded language, but some ambiguous cases remain unresolved.

Corpus Constraints: Several initial aspects (e.g., *Hezbollah*) were discarded because many comments failed the language threshold ($\geq$ 40% English), indicating widespread mixed-language use. Other aspects (e.g., *Zionists*) had very few positive examples, which led to class imbalance that impacted model performance. While augmenting the corpus could mitigate some of these limitations, generating synthetic samples—particularly in sensitive political domains—raises ethical concerns, despite the potential to improve accuracy.

The results on user-generated content from *Deutsche Welle (DW)* and *BBC News* should be interpreted with caution: due to limited short-video content published by these outlets, our corpus contains significantly fewer samples from these sources compared to *TRT World* and *Al Jazeera (AJ)*. This discrepancy does not imply that *DW* or *BBC News* did not cover the aftermath of October 7, but rather reflects platform-specific publishing practices.

Ethics Statement

Our study was conducted in accordance with institutional IRB approval and ethical research standards. We show user-generated content that includes sensitive and potentially harmful language, such expressions that may endorse or glorify violence. These examples are necessary to address our research questions and serve solely for illustrative purposes.

Acknowledgments

This work used Jetstream2 at Indiana University through allocation HUM200003 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

References

Andry Alamsyah, Akmal Wijaya Muharam, and Dian Puteri Ramadhani. 2024. Polarized narratives in digital spaces: A social network examination of the Gaza conflict. In *2024 International Conference on Data Science and Its Applications (ICoDSA)*, pages 527–532.

Carolina. Becatti, Guida Caldarelli, Renaud Lambiotte, and Fabio Saracco. 2019. Extracting significant signal of news consumption from social networks: The case of Twitter in Italian political elections. *Palgrave Commun*, 5(91).

Matthias J. Becker, Laura. Ascone, Matthew Bolton, Alexis Chapelan, Pia Haupeltshofer, Alexa Krugel, Karolina. Placzynta, Marcus Scheiber, and Victor Tschiskale. 2023. Celebrating terror: Antisemitism online after the Hamas attacks on Israel.

Olivia Brown, Laura G. E. Smith, Brittany I. Davidson, Daniel Racek, and Adam Joinson. 2024. Online signals of extremist mobilization. *Personality and Social Psychology Bulletin*.

Ganpat Singh Chauhan and Yogesh Kumar Meena. 2019. Youtube video ranking by aspect-based sentiment analysis on user feedback. In J. Wang, G. R. M. Reddy, V. K. Prasad, and V. S. Reddy, editors, *Soft Computing and Signal Processing*, volume 900, pages 63–71. Springer Singapore.

Nicholas J. Cull. 2008. Public diplomacy: Taxonomies and histories. *The Annals of the American Academy of Political and Social Science*, 616(1):31–54.

Timothy Dozat and Christopher Manning. 2017. Deep biaffine attention for neural dependency parsing. In *5h International Conference on Learning Representations (ICLR 2017)*, Toulon, France.

Joel Finkelstein, Jack Donohue, Alex Goldenberg, Lee Jussim, Alexander Reid Ross, Collin Vasko, Christian Ramos, Taylor Glover, Anisha Jagdeep, Anthony Boyle, and Jiya Modi. 2023. #Twittertakeover: How the Musk acquisition became a launchpad for Gen-Z Neo-Nazis, Ye, and widespread antisemitism. Technical report, Network Contagion Research Institute.

Darina Gold, Marie Bexte, and Torsten Zesch. 2018. Corpus of aspect-based sentiment in political debates. In *Proceedings of the 14th Conference on Natural Language Processing (KONVENS 2018)*, pages 89–99, Vienna, Austria.

Ashley S Jamison, Jason Miklian, Kristian Hoelscher, Witold J Henisz, and Brian Ganson. 2023. Is media sentiment associated with future conflict events? *SSRN Electronic Journal*.

Eduard Kušen and Mark Strembeck. 2023. The effects of multiple exposure to highly emotional social media content during the early stages of the 2022 war in Ukraine. *SN Computer Science*, 4(1):663.

Kristina Lerman, Dan Feldman, Zihao He, and Ashwin Rao. 2024. Affective polarization and dynamics of information spread in online networks. *npj Complexity*, 1(8).

Long Mai and Bac Le. 2021. Joint sentence and aspect-level sentiment analysis of product comments. *Annals of Operations Research*, 300(2):493–513.

Daniel Miehling. 2024. *Online-Antisemitismus verstehen: Hassrede im Web 2.0. Eine vergleichende Studie zwischen der Mainstream-Plattform X und Fringe-Communities auf Telegram*, 1 edition. Nomos Verlag, Baden-Baden.

Kristian Miok, Encarnacion Hidalgo-Tenorio, Petya Osenova, Miguel-Angel Benitez-Castro, and Marko Robnik-Sikonja. 2023. Multi-aspect multilingual and cross-lingual parliamentary speech analysis. *arXiv preprint arXiv:2207.01054*.

Nimra Mughal, Ghulam Mujtaba, Sarang Shaikh, Aveenash Kumar, and Sher Muhammad Daudpota. 2024. Comparative analysis of deep natural networks and large language models for aspect-based sentiment analysis. *IEEE Access*, 12:60943–60959.

Silviu Paun, Ron Artstein, and Massimo Poesio. 2022. *Statistical Methods for Annotation Analysis*. Springer International Publishing, Cham.

Diana Rieger, Lena Frischlich, Stefanie Rack, and Gary Bente. 2020. Digitaler Wandel, Radikalisierungsprozesse und Extremismusprävention im Internet. In *Handbuch Extremismusprävention – Gesamtgesellschaftlich. Phänomenübergreifend*, pages 351–388. Bundeskriminalamt.

Alexander Rietzler, Sebastian Stabinger, Paul Opitz, and Stefan Engl. 2020. Adapt or get left behind: Domain adaptation through BERT language model finetuning for aspect-target sentiment classification. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4933–4941, Marseille, France. European Language Resources Association.

Abuzar Royesh and Shelby Grossman. #COVID1948: How a hashtag on Israel and Palestine went viral [online]. 2021. Freeman Spogli Institute for International Studies, Stanford University.

Eloize Seno, Lucs Silva, Fábio Anno, Fabiano Rocha, and Helena Caseli. 2024. Aspect-based sentiment analysis in comments on political debates in Portuguese: Evaluating the potential of ChatGPT. In *Proceedings of the 16th International Conference on Computational Processing of Portuguese (PROPOR)*, pages 312–320.

Sebastian Stier, Arnim Bleier, Haiko Lietz, and Markus Strohmaier. 2018. Election campaigning on social media: Politicians, audiences, and the mediation of political communication on Facebook and Twitter. *Political Communication*, 35(1):50–74.

Malliga Subramanian, Veerappampalayam Easwaramoorthy Sathiskumar, G. Deepalakshmi, Jaehyuk Cho, and G. Manikandan. 2023. A survey on hate speech detection and sentiment analysis using machine learning and deep learning models. *Alexandria Engineering Journal*, 80:110–121.

Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Liubimov. 2025. Label Studio: Data labeling software. Open source software available from https://github.com/HumanSignal/label-studio.

Lev Topor. 2024. Mis/disinformation and national resilience: Are countries immune to fake news? In *Cyber Sovereignty: International Security, Mass Communication, and the Future of the Internet*, pages 111–131. Springer Nature Switzerland.

Phyu Yi Win Myint, Sin Lo Lo, and Yanchun Zhang. 2024. Unveiling the dynamics of crisis events: Sentiment and emotion analysis via multi-task learning with attention mechanism and subject-based intent prediction. *Information Processing & Management*, 61(4):103413.

Heng Yang, Biqing Zeng, Mayi Xu, and Tianxing Wang. 2021. Back to reality: Leveraging pattern-driven modeling to enable affordable sentiment dependency learning. *CoRR*, abs/2110.08604.

Heng Yang, Chen Zhang, and Ke Li. 2023. PyABSA: A modularized framework for reproducible aspect-based sentiment analysis. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, page 5117–5122, New York, NY, USA. Association for Computing Machinery.

James C. Young, Rudy Arthur, and Hywel T. P. Williams. 2024. CIDER: Context-sensitive polarity measurement for short-form text. *PLOS ONE*, 19(4):e0299490.

Nurulhuda Zainuddin, Ali Selamat, and Roliana Ibrahim. 2016. Improving Twitter aspect-based sentiment analysis using hybrid approach. In *Intelligent Information and Database Systems*, pages 151–160. Springer.

Nurulhuda Zainuddin, Ali Selamat, and Roliana Ibrahim. 2018. Hybrid sentiment classification on Twitter aspect-based sentiment analysis. *Applied Intelligence*, 48(5):1218–1232.

Hao Zhang, Yu-N Cheah, Osamah Mohammed Alyasiri, and Jieuyu An. 2024. Exploring aspect-based sentiment quadruple extraction with implicit aspects, opinions, and ChatGPT: A comprehensive survey. *Artificial Intelligence Review*, 57(2):17.

Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2022. A survey on aspect-based sentiment analysis: Tasks, methods, and challenges (version 2). *arXiv preprint arXiv:2203.01054*.