

Shifting Patterns of Extremist Discourse on Facebook: Analyzing Trends and Developments During the Israel-Hamas Conflict

Rr. Nefriana¹, Muheng Yan¹, Ahmad Diab¹, Wanhao Yu², Deborah L. Wheeler³,
Andrew Miller³, Rebecca Hwa⁴, and Yu-Ru Lin^{1**}

¹ University of Pittsburgh, Pittsburgh, PA 15260, USA

{rr.nefriana, muheng.yan, ahd23, yurulin}@pitt.edu **

² University of North Carolina at Charlotte, Charlotte, NC 28223, USA

wyu6@charlotte.edu

³ United States Naval Academy, Annapolis, MD 21402, USA

{dwheeler, millera}@usna.edu

⁴ George Washington University, Washington, DC 20052, USA

rebecca.hwa@email.gwu.edu

Abstract. This working paper explores trends in extremist Facebook data from July 2023 to June 2024. We examined engagement, sentiment, and topics within Facebook groups categorized as anti-Israel/Semitic, anti-Palestine/Muslim, and anti-both, mapping these trends against five major events related to the recent Israel-Hamas conflict. Our findings support the hypothesis that shifts in trends correspond with these key events, showing varying patterns across different group categories. We observed decreased activity proportion in anti-both groups and increased activity proportion in the two one-sided hate groups at the conflict's onset. This pattern reversed after the Israeli troop withdrawal from Khan Yunis, Gaza. During the conflict, negative content proportion surged, and neutral content proportion fell in all the three group categories. Anti-Palestine/Muslim groups' discourses shifted from religious to social media activism and political/protest around the time the war began, while anti-Israel/Semitic groups moved from political/protest to religious topics a couple of weeks before the war.

Keywords: Israel-Hamas conflict · Online extremism · Hate speech.

1 Introduction

The Israel-Palestine conflict often sparks intense global debates among ethnic groups, political parties, and sovereign states [8]. Social media has extended these discussions beyond geographical limits, enabling global engagement and activism. However, this expansion also amplifies hate speech, particularly against Muslim and Jewish communities [4]. Anti-Muslim and anti-Semitic groups frequently exploit major events [11], with the Hamas-led war in October 2023 fueling increased hatred. This highlights the need for a detailed examination of social media interactions related to anti-Muslim and anti-Semitic content during significant geopolitical events.

Hate speech targeting Muslims and Jews has distinct historical and political roots [2], but the Israel-Palestine conflict intertwines them. Anti-Muslim narratives often exploit Islamophobic sentiments, depicting Muslims as threats to Western values and security [2]. Conversely, anti-Semitic hate usually draws on long-standing conspiracy theories about Jewish influence in politics, economy, and culture [2]. The attacks initiated by Hamas and the ensuing political-military aftermath by Israel appear to align with and intensify these existing prejudices, further polarizing online discourse.

** Corresponding author.

Understanding hate speech and social media engagement during the Israel-Palestine conflict is crucial for several reasons. It helps identify patterns in the rise and fall of online hate speech, offering insights into how major events can either provoke or reduce hatred. This knowledge can inform the development of more effective strategies to counteract online hate by pinpointing key moments for intervention. Additionally, it aids in anticipating potential spikes in hate speech, allowing for timely and proactive measures to mitigate harm.

Despite extensive research on online hate speech [4], there are still gaps in understanding the specifics of hatred during the 2023 Israel-Palestine war. First, the variation in hate speech engagement throughout the conflict needs more investigation. It is important to determine if spikes in engagement align with specific events—military, political, or diplomatic—and to analyze the patterns of these fluctuations. Second, the link between major events in the conflict and changes in online sentiments within anti-Muslim and anti-Jewish groups is underexplored. While previous studies indicate that sentiments can shift rapidly in response to news and events [5], the direction and correlation of these changes with specific incidents have not been studied in the context of mutual hatred in a conflict. Lastly, examining how discussion topics on social media evolve during the conflict can provide additional insights. Identifying shifts in focus—such as reports on attacks, religious debates, empathy for war victims, or condemnation of perpetrators—can help understand how hate groups sustain and expand their influence.

To address this gap, this study proposes the following research questions:

- **RQ1:** How does engagement with hate speech vary as the Israel-Palestine conflict develops? Are these changes related to major events?
- **RQ2:** How do sentiments within anti-Muslim and anti-Jew groups change during the conflict? Are these changes related to major events?
- **RQ3:** How do the topics of social media discussions change within these groups over time?

This research provides an understanding of the dynamic of online hate speech during the Israel-Palestine conflict, which will contribute to the broader discussion on online hatred to promote a more inclusive digital space.

2 Related Works

Hate speech on social media is a pervasive issue, amplified by the reach and anonymity these platforms offer. Online hate speech, which disparages individuals based on characteristics like race, religion, or gender, is facilitated by the unregulated nature of social media [16]. The prevalence is alarmingly high, with 80% of people in the EU encountering it and 40% feeling personally attacked or threatened [7]. This contributes to an environment of prejudice and intolerance, which can escalate to violence [12]. Platforms like Facebook and Twitter play a key role in mediating and amplifying hate speech [4]. Their design and policies affect the frequency and spread of hate speech, with anonymity features protecting harassers and algorithmic suggestions promoting racist content [1].

Hate speech online can manifest as religious hate speech, racism, political hate, and gendered hate [4]. For example, anti-Islam hate, a focus of this study, is prevalent on social media and often associated with Islamophobia and global cultural processes [10]. Violent events, such as terrorist attacks, are linked to surges in online hate speech; for instance, the #StopIslam hashtag trended on Twitter after the Brussels attacks in 2016, spreading racialized hate speech against Muslims [13]. Geo-political events also impact online hate speech, with the Israeli-Palestinian conflict being a key example. During intensified conflicts, like military operations in Gaza, there is a noticeable increase in anti-Semitic content. This pattern is consistent in various European countries, where anti-Semitism

and anti-Zionism often overlap, complicating the distinction between political criticism and outright hate [2]. The escalation of hostilities in the Israel-Hamas conflict contributes to the spread of hate speech, with online platforms filled with content reinforcing existing prejudices and stereotypes [2].

Given these dynamics from past research, we build an understanding of how such conflict-driven hate speech would evolve. This study aims to fill the gaps in existing research by examining the trends in engagement, sentiment, and topics within extremist Facebook groups throughout the Israel-Hamas conflict.

3 Methodology

3.1 Data

We examined Facebook data to answer the research questions. To create a corpus of hate speech against Jews and Muslims, we start with a list of keywords curated by an expert on Middle Eastern regional politics. This list contains phrases and slogans that explicitly express hate against either Muslim or Jewish groups. This list includes explicit phrases like "death to Israel" and "Arabs are inferior," as well as conflict-related terms such as "From the Rivers to the Sea"⁵ and "Palestinian People do not Exist."

We used the CrowdTangle API to identify and collect Facebook group posts containing these keywords. The search was conducted in October and November 2023, after Hamas launched rockets into Israel on October 7, 2023. Three graduate students reviewed the groups and posts to determine if the posts expressed hate and if the groups focused on Middle Eastern or Israel-Palestine issues. They identified 128 groups potentially engaged in hate speech. One author reviewed and annotated the results, leading to an agreement on 32 groups. We labeled 10 groups for anti-Muslim speech, 18 for anti-Semitism, and 4 for both. We collected historical posts from these groups from July 2023 to June 2024. These 32 groups produced 1,284,331 posts in eleven months. Each post is labeled by its group: "anti-Palestine/Muslim," "anti-Israel/Semitism," "anti-both," or "other."

3.2 Sentiment Analysis

After removing stop words from each Facebook post in our dataset, we conducted sentiment analysis using a pre-trained model from Hugging Face, `finiteautomata/bertweet-base-sentiment-analysis` [15]. This model uses BERTweet as the base model and was trained on the SemEval 2017 corpus. Our sentiment analysis resulted in each post being tagged with either positive, neutral, or negative sentiment.

3.3 Topic Modeling

We modeled the Facebook corpus topics using BERTopic [9] models. The 160,456 English posts are tokenized into sentences for training the BERTopic model. We initialized the semantic module with pretrained sentence BERT and update the class-based TF-IDF on our sentence-level Facebook corpus [9]. Topic clusters were identified using HDBSCAN [3] with a minimum cluster size of 1,000 sentences. This process yielded 263 fine-grained topic clusters, which are aggregated into six major clusters based on their semantic centroids and top keywords, as detailed below.

- 1) **Social Media Activism:** Includes "call for action" phrases and slogans, such as "ImWithIsrael."

⁵ Used by both anti-Muslim and anti-Semitic groups to express the desire to eliminate Jews/Muslims completely from the Jordan River to the Mediterranean Sea

- 2) **Religious**: Contains references or metaphors from religious texts like the Quran, Bible, and Tanakh.
- 3) **Politics and Protests**: Involves mentions of Western political figures and information about offline protests.
- 4) **Military**: Covers reports and discussions on military actions in Gaza, such as bombings or ground assaults.
- 5) **Ethnicity-based Hate and War-crime Accusations**: Includes explicit hate towards Muslim or Jewish groups and accusations of war crimes in the conflict.
- 6) **Irrelevant Topics**: Contains texts unrelated to the Israel-Palestine conflict, such as social media greetings and advertisements.

We applied the trained BERTopic model to the 160,456 posts, creating a topic distribution matrix where each post is represented by a 5-dimensional topic vector (excluding the irrelevant topic).

3.4 Major event selection

We selected key events for our analysis based on the highlighted events from the Wikipedia’s page "Timeline of the Israel– Hamas war" ⁶. These major events include the initial attack, the invasion of the Gaza Strip, and the first ceasefire. There are two other major events on that list: the Yemen airstrikes and the Rafah offensive. In our exploration, we noticed that the first three major events often coincided with significant changes in the trends of engagement, sentiment, and/or topic, and hence we continue using these three events in our analysis. However, since we rarely observed major changes in our time series regarding the Yemen airstrikes, we omitted this event from our current analysis to increase clarity. Furthermore, we skipped the Rafah offensive that started on May 7, 2024, because we missed 10 days of data following this event, hindering our ability to determine trends around that period. Instead, we added two major events that often coincided with significant changes in our time series: the major attack on Rafah [14] and the withdrawal of Israeli troops from Khan Yunis [6]. The list of these major events with the corresponding dates is available in Table 1.

Table 1: Key events related to the recent Israel-Hamas we used in our analysis.

No.	Date	Event
1.	October 7, 2023	Hamas attacked Israel.
2.	October 28, 2023	Israel started the invasion of the Gaza Strip.
3.	November 24, 2023	The seven-day ceasefire formally started.
4.	December 15, 2023	Israel initiated the first major assault on Rafah (Gaza).
5.	April 7, 2024	Israeli troops withdrew from Khan Yunis (Gaza).

4 Analysis

4.1 Trends in Social Media Engagement.

Initial Engagement Dynamics and Conflict Onset. From July 1, 2023, to June 30, 2024, the engagement dynamics in Facebook groups showed notable fluctuations, as depicted in Figure 1. Figure 1A highlights the distribution of post proportions across different group categories over time. While the proportions of posts across the three group

⁶ https://en.wikipedia.org/wiki/Timeline_of_the_Israel-Hamas_war

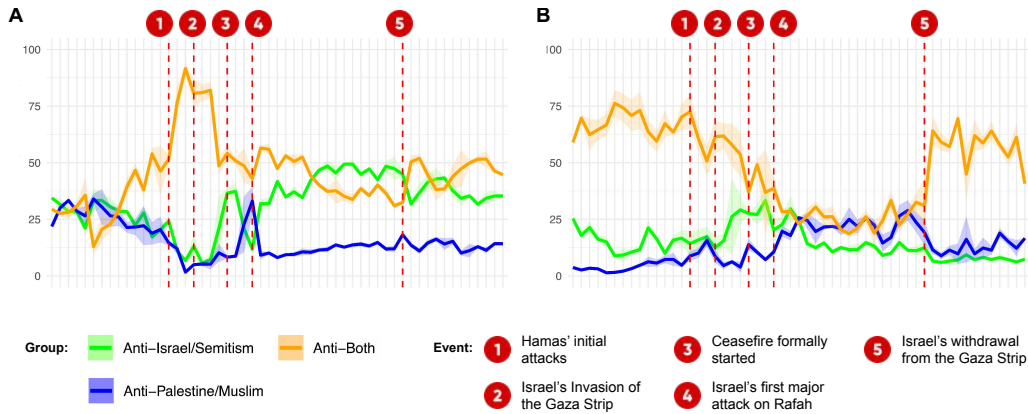


Fig. 1: The engagement of Facebook groups from July 1st, 2023 to June 30, 2024: (A) the proportions of posts over time, and (B) the proportions of interactions per post over time. The data are aggregated weekly. Solid lines are medians, while shaded areas are 95% CI. Statistics are normalized to percentages for anti-Israel/Semitism, anti-Palestine/Muslim, anti-both, and anti-others. The anti-others category is omitted for clarity.

categories were similar around mid-2023, a few weeks before Hamas's attack on October 7th, the proportion of posts from the anti-both groups started to increase steeply. At the start of the latest war, the proportion of posts from the anti-both groups was the highest. However, the proportion of posts from the anti-both groups decreased significantly with the onset of the conflict, marked by Hamas's initial attack and Israel's subsequent invasion of Gaza.

This gradual decrease in the proportion of posts from the anti-both group was accompanied by a significant rise in the proportion of posts from anti-Israel/anti-Semitism groups. Albeit not as sharply as that from anti-Israel/anti-Semitism groups, the proportion of posts from anti-Palestine/anti-Muslim groups also gradually rose. It is possible that, with the onset of the war, individuals focused their attention more on one side than on both sides. Future research should explore these phenomena further to test this hypothesis.

Some significant changes were also observed in the percentage of interactions received per post (Figure 1B). A few weeks before the war, posts from the anti-both group category received the highest interactions, outpacing those from the anti-Israel/anti-Semitism and anti-Palestine/Muslim categories. However, after the onset of the war, the proportion of posts from the anti-both groups decreased sharply, while those from the anti-Israel/anti-Semitism and anti-Palestine/Muslim groups generally increased, despite some fluctuations.

Another interesting change during the war was the reversal of interaction patterns: before the war, anti-Israel/Semitism groups had a higher interaction per post proportion compared to anti-Palestine/Muslim groups. This pattern reversed after the onset of the war, particularly after the Rafah attack, with posts from anti-Palestine/Muslim groups receiving more interactions than those from anti-Israel/Semitism groups. This pattern in interaction proportions during the war contrasts with the pattern in post proportions: the proportion of posts from anti-Israel/Semitism groups was higher than that of anti-Palestine/Muslim groups. This may indicate that during the war, users in anti-Israel/Semitism groups were more focused on creating posts than those in anti-Palestine/Muslim groups, but the opposite was true in terms of generating interactions.

Reversals in Engagement Trends Around Israel's Troop Withdrawal. After the start of the war, the proportion of posts from anti-Israel/Semitic groups was the

highest. However, around the time of Israel’s withdrawal from Khan Yunis, the pattern reversed, with a slight decrease in the proportion of posts from anti-Israel/Semitic groups and a slight increase in the proportion of posts from the anti-both groups.

Meanwhile, during the war, the interactions per post among the anti-both groups and anti-Palestine/Muslim groups became nearly equal, at around 25%. The proportion of interactions per post from the anti-Israel/Semitic groups was slightly lower than those of the anti-both and anti-Palestine/Muslim groups. This distribution quickly changed around Israel’s withdrawal from Khan Yunis. At this point, the interactions per post among the anti-both groups rose significantly, while the proportions from the anti-Palestine/Muslim and anti-Israel/anti-Semitic groups decreased. These changes during Israel’s troop withdrawal reflect patterns opposite to those observed at the onset of the war.

The Post Rates in Different Groups During the Ceasefire. The ceasefire that began on November 24 was associated with a decline in the percentage of posts from anti-both groups. At the same time, the growth rate of posts from anti-Israel/Semitism groups nearly stopped and significantly reduced one week later. The ceasefire was also followed by an increase in posts from anti-Palestine/anti-Muslim groups a week later. This pattern may suggest that the ceasefire, which temporarily halted Israel’s invasion, reduced engagement proportion in anti-both groups and prevented a surge in negative views toward Israel/Jews, leading to an increase in negative view proportion toward Palestine/Muslims.

Variation in Group Posts During Major Military Actions. Israel’s major attack on Rafah on December 15 was associated with a notable shift: the proportion of posts from anti-both groups increased by approximately 12.5%. Post proportion from anti-Israel/Semitism groups also saw a rise of almost 20%, while those from anti-Palestine/Muslim groups decreased by 25%. Previously, around the time of Israel’s first attack on Gaza on October 28, we observed a marginal increase in anti-both group posts as well. However, this slight increase was accompanied by a marginal decrease in posts from anti-Israel/anti-Semitism groups. We conjecture that the first attack on Gaza was perceived as more justified due to its nature as a response to a Hamas attack, resulting in a slight decrease in the proportion of posts in anti-Israel/anti-Semitism groups compared to the significant increase observed when Israel launched its first major attack on Rafah.

4.2 Sentiment Analysis Across Different Facebook Groups

Anti-Israel/Semitism Groups: Predominantly Negative Sentiments. Figure 2A illustrates the sentiments over time within anti-Israel/Semitism groups. These groups generally exhibited less neutral and more negative sentiments compared to anti-Palestine/Muslim groups (Figure 2B). The onset of the conflict, marked by Hamas’s attack and Israel’s responses, coincided with an increase in negative sentiment and a decrease in neutral sentiments within both group categories.

Anti-Both Groups: Drastic Shifts in Sentiment. Figure 2C displays the sentiment trends in anti-both groups, which experienced more dramatic changes compared to the other two categories. These groups were more neutral before the conflict than other groups, but the neutrality significantly decreased during the war, accompanied by a sharp increase in negativity. Following the withdrawal of Israeli forces from Khan Yunis, neutrality increased while negativity decreased.

Major Events and Sentiment Trends. Significant events coincided with these trends moving in opposite directions. For example, Israel’s invasion of Gaza on October 28 corresponded with a rise in negative sentiments in the anti-Israel/Semitism groups and a de-

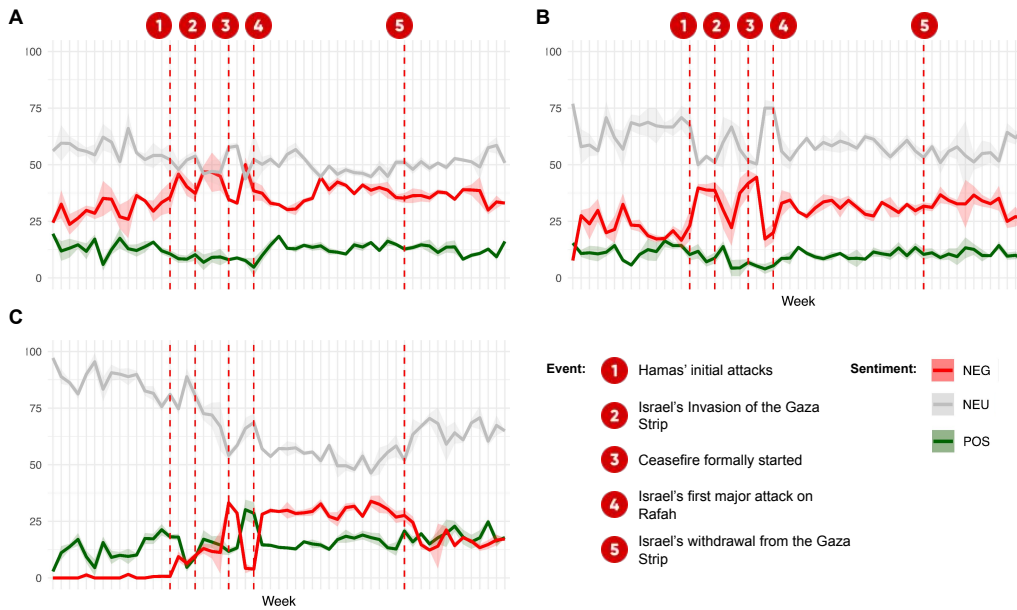


Fig. 2: The percentages of posts with their corresponding sentiment on the Facebook groups over time (from July 1st, 2023, to June 30, 2024): (A) from anti-Israel/Semitic groups, (B) from anti-Palestine/Muslim groups, and (C) from anti-both groups. The data are aggregated weekly. The solid lines represent the medians, while the shaded areas represent the 95% CI.

crease in negative sentiments in the anti-Palestine/Muslim groups. Differently, the ceasefire was associated with a slight decrease in negative sentiments within anti-Israel/Semitism groups and a slight increasing trend in negative sentiments within anti-Palestine/Muslim groups.

Regarding the anti-both groups, the proportion of positive sentiment was higher than that of negative sentiment before the war. However, this pattern reversed during the war. By the end of our study period, the proportion of positive sentiment had marginally surpassed negativity once again, albeit by a very small difference.

4.3 Topic Trends Across Different Facebook Groups

Figure 3 breaks down the trend of daily topic probabilities across various Facebook group categories and overall from July 1, 2023, to June 30, 2024.

Anti-Israel/Semitism Groups: Social Media Activism and Religious Topics.

In the anti-Israel/Semitism groups (Figure 3A), in general, the most prominent topic discussed was social media activism, followed by religious topics. Before the conflict began, these groups focused more on political and protest topics than religious ones. However, this trend shifted approximately two weeks before the Hamas attack on October 7, with religious topics becoming more prevalent than political discussions.

Anti-Palestine/Muslim Groups: Predominance of Religious Discussions.

In contrast, the anti-Palestine/Muslim groups (Figure 3B) predominantly discussed religious topics, especially before the first major attack on Rafah. Social media activism was the next most discussed topic. After the first major attack on Rafah, there was a noticeable decline in religious discussions within these groups, while social media activism and political/protest topics gained a little more prominence.




Fig. 3: The distribution of topics on the Facebook groups over time (from July 1st, 2023, to June 30, 2024): (A) in anti-Israel/Semitic groups, (B) in anti-Palestine/Muslim groups, and (C) in anti-both groups. The data are aggregated weekly. The solid lines represent the medians, while the shaded areas (which are not too apparent due to the small daily variability) represent the 95% CI).

Anti-Both Groups: Shifts in Discussion Topics. Among the anti-both groups (Figure 3C), religious and social media activism topics were discussed the most and at similar rates overall. However, religious topics were slightly more prevalent before the first major attack on Rafah. Between the attack on Rafah and the withdrawal of Israel’s troops from Gaza, social media activism topics were discussed more than religious ones. After the Israel’s troops withdrawal, the probabilities of the two topics became very similar. Similar to anti-Palestine/Muslim groups, political/protest topics ranked third as the most prominent topics discussed in the anti-both groups and began to be discussed more frequently when the war started.

4.4 Qualitative Answers for the RQs

The analysis conducted in this study provides several insights into the research questions posed. Regarding **RQ1** (How does engagement with hate speech vary as the Israel-Palestine conflict develops?), the data indicates that engagement levels fluctuated significantly in response to key events. Notably, the onset of the conflict was associated with an increase in single-sided posts (either anti-Israel/anti-Semitism or anti-Palestine/Muslim), suggesting a polarization in user engagement. For **RQ2** (How do sentiments within anti-Muslim and anti-Jew groups change during the conflict?), the data reveals that both groups exhibited a marked increase in negative sentiments coinciding with major conflict events, such as Hamas’s initial attack and Israel’s invasion of Gaza. However, the ceasefire period saw a complex shift, with a decrease in negative sentiments within anti-Israel/anti-Semitism groups and an increase in anti-Palestine/Muslim groups. Lastly, for **RQ3** (How do the topics of social media discussions change within these groups over time?), the analysis highlights that the dominant topics within these groups shifted notably during the conflict. Anti-Israel/Semitism groups transitioned from discussions on political/protest topics to religious ones, although this change started before the war.

Anti-Palestine/Muslim groups initially concentrated on religious topics, which decreased after the first major attack on Rafah, giving way to more social media activism and political/protest discussions. These findings show the dynamics of online hate speech targeting Muslims and Jews, and the relationships between geopolitical events and the narratives within groups.

5 Discussion and Conclusion

In this study, we examined the dynamics of Facebook groups from July 1, 2023, to June 30, 2024, focusing on engagement, sentiments, and topics discussed. We classified these groups into four categories: anti-Israel/Semitism, anti-Palestine/Muslim, anti-both, and other. We then analyzed how trends in the first three categories corresponded with five major events of the Israel-Hamas war.

Our findings support the hypothesis that Facebook group dynamics coincided with major events. The onset of the conflict, marked by Hamas's attack and subsequent Israeli actions, coincided with a decrease in content creation proportion by anti-both users and an increase among anti-Israel/Semitism and anti-Palestine/Muslim users. This pattern reversed after the Israeli troop withdrawal from Gaza, with activity proportion rising in the anti-both groups and falling in the other two categories. During the war's start, negative content percentage increased, and neutral content percentage decreased in the all group categories. Anti-Palestine/Muslim groups shifted from religious to social media activism and political/protest discourse, while anti-Israel/Semitism groups moved from political/protest to religious topics, although this shift began before the war.

These findings have several implications. The coinciding changes in engagement, sentiment, and topic trends suggest a potential correlation between online and offline expressions of hate. This provides a basis for a more refined time series analysis in the future. The Facebook groups exhibited more negativity and less neutrality during the war, indicating a shift in discourse patterns. It suggests that the social media platforms may implement effective strategies to combat negativity without impeding constructive criticism, while it poses challenges for community moderators to carefully differentiate between critical feedback and hate speech.

5.1 Limitation and Future Works

This study has several limitations, which also point to directions for future research. The dataset covers only one year, providing a limited view of the conflict's complex and evolving nature. The qualitative approach, focusing on event coincidences, restricts our understanding of causal relationships. Future research should combine quantitative and qualitative methods to explore these relationships more thoroughly. Additionally, while this analysis grouped anti-Israel with anti-Semitism and anti-Palestine with anti-Muslim views for simplicity, future studies should aim to distinguish between these concepts more precisely to better understand the dynamics of hostility toward these groups. Lastly, regarding the methodological consideration, the graduate students and the reviewer who examined the groups and posts to assess whether they expressed hate speech were not from either country involved in the conflict. While this likely helped reduce potential biases in their review, we recognize that personal biases, if any, can never be entirely eliminated, which remains a limitation of the study.

Disclosure of Interests. The authors declare no conflicts of interest. While it is acknowledged that personal backgrounds may have some potential for influence on the research, the analysis was conducted with a focus on impartiality.

Acknowledgement The authors would like to acknowledge support from AFOSR, ONR, Minerva, NSF #2318461, and Pitt Cyber Institute’s PCAG awards. The research was partly supported by Pitt’s CRC resources. Any opinions, findings, and conclusions or recommendations expressed in this material do not necessarily reflect the views of the funding sources.

References

- [1] Ben-David, A., Fernández, A.M.: Hate speech and covert discrimination on social media: Monitoring the facebook pages of extreme-right political parties in spain. *International Journal of Communication* **10**, 27 (2016)
- [2] Berez, T., Devinat, C.: Relevance of cyber hate in europe and current topics that shape online hate speech. Retrieved October 7, 2020 (2017)
- [3] Campello, R.J., Moulavi, D., Sander, J.: Density-based clustering based on hierarchical density estimates. In: *Pacific-Asia conference on knowledge discovery and data mining*. pp. 160–172. Springer (2013)
- [4] Castaño-Pulgarín, S.A., Suárez-Betancur, N., Vega, L.M.T., López, H.M.H.: Internet, social media and online hate speech. systematic review. *Aggression and violent behavior* **58**, 101608 (2021)
- [5] Chetty, N., Alathur, S.: Trigger event and hate content: Insights from twitter analytics. In: *2019 International Conference on Advances in Computing and Communication Engineering (ICACCE)*. pp. 1–5. IEEE (2019)
- [6] CNN: April 7, 2024 - israel-amas war (2024), <https://www.cnn.com/middleeast/live-news/israel-amas-war-gaza-news-04-07-24/index.html>, accessed: 2024-07-19
- [7] Gagliardone, I., Gal, D., Alves, T., Martinez, G.: *Countering online hate speech*. Unesco Publishing (2015)
- [8] Gelvin, J.L.: *The Israel-Palestine conflict: One hundred years of war*. Cambridge University Press (2014)
- [9] Grootendorst, M.: Bertopic: Neural topic modeling with a class-based tf-idf procedure. arXiv preprint arXiv:2203.05794 (2022)
- [10] Horsti, K.: Digital islamophobia: The swedish woman as a figure of pure and dangerous whiteness. *New Media & Society* **19**(9), 1440–1457 (2017)
- [11] Krämer, G.: Anti-semitism in the muslim world: A critical review. *Die Welt des Islams* **46**(3), 243–276 (2006)
- [12] Papacharissi, Z.: Democracy online: Civility, politeness, and the democratic potential of online political discussion groups. *New media & society* **6**(2), 259–283 (2004)
- [13] Poole, E., Giraud, E., de Quincey, E.: Contesting# stopislam: The dynamics of a counter-narrative against right-wing populism. *Open Library of Humanities* **5**(1) (2019)
- [14] Post, T.J.: Idf launches major attack on rafah, philadelphia corridor in gaza’s south. *The Jerusalem Post* (2023), <https://www.jpost.com/israel-amas-war/article-778165>, accessed: 2024-07-19
- [15] Pérez, J.M., Giudici, J.C., Luque, F.: pysentimiento: A python toolkit for sentiment analysis and socialnlp tasks (2021)
- [16] Zhang, Z., Luo, L.: Hate speech detection: A solved problem? the challenging case of long tail on twitter. *Semantic Web* **10**(5), 925–945 (2019)