Investigating Polarization in YouTube Comments via Aspect-Based Sentiment Analysis

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

Abstract

We investigate the use of Aspect-Based Sentiment Analysis (ABSA) to analyze polarization in online discourse. For the analysis, we use a corpus of over 3 million user comments and replies from four state-funded media channels from YouTube Shorts in the context of the 2023 Israel–Hamas war. We first annotate a subsample of approx. 5 000 comments for positive, negative, and neutral sentiment towards a list of topic related aspects. After training an ABSA model (Yang et al., 2023) on the corpus, we evaluate its performance on this task intrinsically, before evaluating the usability of the automatic analysis of the whole corpus for analyzing polarization. Our results show that the ABSA model achieves an F1 score of 77.9. The longitudinal and outlet analyses corroborate known trends and offer subject experts more fine-grained information about the use of domain-specific language in user-generated content.

1 Introduction

Polarization in online discourse has increased over the last decade. However, polarization is a linguistic phenomenon that is difficult to define in a way that allows for an automatic analysis of large corpora. Prior work has investigated polarization primarily through small-scale or short-term datasets (Alamsyah et al., 2024; Zeitzoff, 2016), primarily qualitative analyses (Harel et al., 2020), or coarse computational methods such as stance detection or user clustering, often applied to Twitter data (Becatti et al., 2019). This focus reflects the dominant role of Facebook and Twitter as platforms for political communication and information-seeking (Stier et al., 2018), which has led to a research landscape heavily concentrated on these two platforms. In contrast, polarization processes are rarely examined on YouTube or using longitudinal datasets that incorporate fine-grained sentiment analysis focused on specific actors or ideological targets.

One way of approaching the automatic analysis of polarization is to investigate sentiment towards central actors and topics in the discourse in question. We propose to use Aspect-Based Sentiment Analysis (ABSA) to analyze such polarization. ABSA is particularly suited to analyzing complex opinions in online discourse, where a single sentence may express multiple sentiments toward different entities (Chauhan and Meena, 2019; Mai and Le, 2021). Traditional sentence-level sentiment analysis is limited in such contexts, as it assigns a single sentiment label to an entire sentence. In contrast, ABSA allows for a more granular analysis, enabling targeted sentiment detection for each mentioned entity. The following example illustrates multiple sentiments directed toward different aspects:

(1) *So what. The Israeli government IS committing genocide. Just look at the civilian death tolls. That doesn't diminish the need to condemn Hamas for their terrorist and bloody attack on civilians. Both can be true at the same time. Greta is awesome. She didn't back down, but she remained calm. As a fellow Swede, I'm very proud of her.*

This comment illustrates how sentiment is distributed across multiple entities: *Israel*, *Hamas*, *Greta*, *Swede*. Each of those is associated with distinct emotional evaluations. ABSA enables the detection of such distinctions, providing insight into how users express affective positions toward various actors within the same comment. When applied to large-scale discourse, this fine-grained mapping of sentiment allows us to observe how alignment with or opposition to ideological actors shifts over time, offering a complementary linguis-

Proceedings of Recent Advances in Natural Language Processing, pages 718–728
Varna, Sep 8–10, 2025
https://doi.org/10.26615/978-954-452-098-4-083

tic lens to polarization as a dynamic process driven by social influence and persuasive communication (Treuillier et al., 2024).

Our data consists of user-generated content tied to YouTube Shorts, vertical videos, often under 60 seconds in length, closely resembling the content format of TikTok. We have collected over 3 million comments and replies across four state-funded media channels (Al Jazeera, TRT World, BBC News, and Deutsche Welle) to examine emotional responses toward geopolitical actors (Israel, Hamas, Palestinians, Zionists, Jews, and Muslims) over a 12-month period of the Israel-Hamas war[1]. Our goal for the work presented here is to determine whether we can use a current ABSA model (Yang et al., 2021, 2023) to gain insights into polarization. This requires an intrinsic analysis of the quality of the ABSA system on the current task along with an extrinsic evaluation of the usability of ABSA for analyzing polarization.

This study applies Aspect-Based Sentiment Analysis (ABSA) to model fine-grained sentiment trajectories towards geopolitical actors. If the approach is successful, it will enable us to trace how affective alignments shift in response to distinct conflict-related events across state-funded media contexts.

Our work addresses the following research questions:

RQ1: How well is a current ABSA model suited to analyze sentiment towards geopolitical actors?

RQ2: How well does the automatic sentiment analysis model known trends in polarization?

The remainder of this paper is structured as follows: We describe related work in Section 2 and our methodology in Section 3. Section 4 discusses the results of RQ1, and Section 5 the results of RQ2. We conclude in Section 6.

2 Related Work

Prior research has extensively addressed hate speech, misinformation, bias, and political extremism in online spaces (Brown et al., 2024; Finkelstein et al., 2023; Rieger et al., 2020), including in the context of the Israel-Hamas War (Becker et al., 2023; Miner and Ortega, 2024; Zaghouani et al., 2024; Topor, 2024). Building on these foundations, sentiment analysis has emerged as a key method to model ideological alignment, polarization, and user engagement in socio-political discourse. Studies such as the one by Jamison et al. (2023) demonstrate that shifts in media sentiment can precede conflict outbreaks, while others have shown how sentiment dynamics reflect affective polarization (Lerman et al., 2024).

Existing studies have applied sentiment and stance analysis to online political discourse, e.g., in the contexts of US Supreme Court decisions (Wang et al., 2014), Brexit (Vorakit et al., 2020), gun control and immigration (Roy and Goldwasser, 2021), political debate on Twitter (Sirisha and Chandana, 2022), and various forms of longitudinal polarization (Amendola et al., 2024).

Recent research further expands the methodological scope of sentiment modeling in socio-political event analysis. Muñoz et al. (2024) proposed a sentiment-informed algorithm to track political polarization across online networks. Kušen and Strembeck (2023) demonstrated how emotional exposure on social media during the early stages of the Ukraine war influenced user affect over time, while Win Myint et al. (2024) advanced multi-task emotion classification in crisis discourse, identifying base emotions such as anger, fear, and sadness. 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.

ABSA has gained traction as a method for extracting sentiment toward discrete entities or themes, particularly in domains such as hate speech detection (Mughal et al., 2024; Zainuddin et al., 2016, 2018; Zhang et al., 2024), political communication (Gold et al., 2018; Miok et al., 2023; Seno et al., 2024), and changing attitudes towards areas of a country (Davoodi et al., 2024). ABSA offers more scalable solutions for user-generated discourse, enabling robust, transformer-based sentiment extraction at the aspect level across multiple targets within the same comment (Rietzler et al., 2020; Zhang et al., 2022), making it suitable for modeling longitudinal sentiment variation in politically charged online discussions.

[1]Note that we are not taking a political stance on this conflict, our interest is in analyzing how different sides polarize.

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: Statistics of the full dataset across media outlets.

3 Methodology

3.1 Dataset

We collected the dataset from YouTube via the official YouTube API[2]. Data collection began approximately two weeks after the Hamas-led attacks on Israel on October 7, 2023[3]. To narrow the scope, we focused on collecting posts in English from four state-funded media outlets: *TRT World* (Turkey), *BBC News* (United Kingdom), *Al Jazeera English* (Qatar), and *Deutsche Welle* (Germany)[4]. These outlets represent divergent geopolitical perspectives and serve as influential platforms for shaping regional and global narratives. We collected all Shorts and manually selected a subset based on their topical relevance. Over a twelve-month data collection period, we gathered more than 3.4 million comments and replies associated with these four sources, generated by 787 157 unique users. Table 1 provides an overview of the number of unique video IDs, user-generated comments, and distinct user counts per outlet.

3.2 Preprocessing

A multi-step preprocessing pipeline was implemented to prepare the dataset for analysis. Files were scanned for structural validity, filtering malformed JSON lines and removing entries with missing fields. User mentions (e.g., `@username`) were normalized, and only English-language comments were retained using the `langdetect` library[5]. To ensure linguistic quality, we computed an English lexical coverage ratio based on the `nltk.corpus.words` vocabulary, discarding comments with less than 40% English terms[6].

After preprocessing, $n = 868\,094$ comments and replies remained from an initial $n = 1\,071\,630$ aspect-bearing instances.

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 2: Distribution of annotated segments by aspect and sentiment label.

3.3 Annotating Aspect-Based Sentiment

The annotated dataset comprises 5 147 text segments / sentences[7]. We selected sentences containing the following aspect terms: *Jews*, *Israel*, *Zionists*, *Palestinians*, *Palestine*, *Hezbollah*, and *Hamas*. Due to a low number of occurrences, *Hezbollah* was excluded from the data. To ensure linguistic diversity, we discarded near-duplicate or semantically redundant content and selected varied comments and replies per aspect category. The selection of segments was carried out iteratively with the goal of balancing the data with respect to sentiment as far as possible in order to avoid prejudicing the classifier towards a specific sentiment. However, this was not always possible: Especially the aspect term *Zionists* predominantly occurs with negative sentiment, with only a limited number of positive samples. Table 2 summarizes the final distribution of annotated examples by aspect and sentiment.

We performed the manual annotation of the selected data using `Label Studio`, an open-source data labeling tool (Tkachenko et al., 2025). The annotations combine identifying spans containing

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

[3] A retrospective attempt to collect pre-October 7 data revealed that nearly 50% of the original material had been deleted, likely due to YouTube's evolving content moderation policies.

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

[5] https://pypi.org/project/langdetect/

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

[7] Each segment corresponds to a dependency-parsed analysis containing at least one target aspect. While most segments are sentence-length, some extend across multiple clauses or sentences, resulting from informal punctuation, ellipses, and unstructured phrasing in user-generated content, which occasionally led the parser to treat longer passages as single units.

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 3: Inter-annotator agreement scores across sentiment classes.

Aspect	Train	Test
Hamas	996	227
Hezbollah	24	7
Israel	770	180
Jews	443	102
Muslims	494	124
Palestine	348	114
Palestinians	640	153
Zionists	406	119

Table 4: Train/test split of annotated dataset by aspect category.

relevant aspect categories and sentiment labels. Aspect spans were chosen based on dependency parses (see Section 3.4). Sentiment was annotated as *positive*, *neutral*, or *negative*. The annotations were primarily conducted by a domain expert with specialized knowledge in political communication, linguistics, and online discourse, including expertise in coded language and sentiment analysis.

To assess annotation reliability, we followed the approach of Agarwal et al. (2011) and constructed a stratified sample of 500 annotations, to ensure balanced coverage of aspect categories and sentiment labels. From the (*aspect*, *sentiment*) pairs, we uniformly sampled up to a fixed per-group quota. Where individual categories contained insufficient data, sampling with replacement was applied to preserve proportionality. This sample was independently labeled by a second expert and was used to calculate inter-annotator agreement (IAA). We report IAA in Table 3, showing a strong inter-annotator agreement across all sentiment classes, with Cohen's κ = 0.862 and Krippendorff's α = 0.855.

To create the training and test data for the ABSA model, we performed an 80/20 random split of the annotated dataset. See Table 4 for statistics of the training and test data.

We used the test set for the intrinsic evaluation before running the model on the full dataset for the extrinsic evaluation.

3.4 Dependency Parsing

We used the biaffine graph-based dependency parser (Dozat and Manning, 2017), implemented in the `SuPar` library[8], to extract syntactic structures and align aspect terms with their grammatical heads. Parsing was performed using the pretrained `biaffine-dep-en` model, trained on English Universal Dependencies (UD) treebanks. The parser was run on a GPU using batch-wise prediction over tokenized user-generated content. While dependency parsing enhanced aspect-term extraction, it did not fully resolve ambiguities in sentiment attribution. Future improvements may result from incorporating contrastive learning techniques or domain-adapted embeddings to better capture context-sensitive sentiment in highly polarized political discourse.

3.5 ABSA Model

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). We initially compared Microsoft's `DeBERTa-v3-base` (He et al., 2021) with `DeBERTa-v3-large-absa-v1.1`. The latter outperformed the base model on a held-out validation subset of our data and was therefore selected for task-specific finetuning. The pretrained model was trained on English-language benchmark datasets including Twitter and SemEval corpora for aspect-based sentiment classification[9]. We then finetuned it for 5 epochs on our task-specific training set using an NVIDIA A100 GPU. We set the batch size to 2 and applied gradient accumulation over 8 steps (effective batch size of 16). We used a cosine learning rate scheduler with an initial learning rate of 1×10^{-5}. Since we did not define a separate development set, we selected the best model checkpoint based on validation loss. To align with the aspect-prompted classification (APC) framework, aspect terms in each sentence were replaced with a `$T$` marker. Sequences were truncated at 512 tokens.

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

[9] While the model was originally trained on social media data, our domain is considerably different from the training data, especially wrt. the specific targets we are interested in. Performance on a held-out subset indicated the need for further domain adaptation via finetuning.

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.

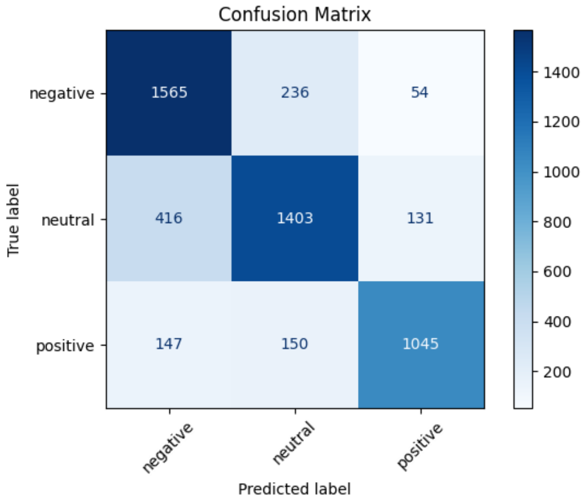

Figure 1: Confusion matrix for sentiment classification on the test set.

4 RQ1: Model Evaluation on Test Data

To evaluate model performance intrinsically, Table 5 reports the overall precision, recall, and across sentiment classes. The model achieves an overall F1 of 77.9, demonstrating robust performance in classifying explicit sentiment expressions for our aspect terms. Figure 1 shows the confusion matrix of the experiment. This shows that the model most frequently misclassifies neutral sentences as negative, the opposite direction causes the second highest error rate.

Since we are interested in the sentiment towards specific targets, we also evaluated how well the model performs on individual aspects. Table 6 shows these results. While overall performance is robust, notable differences exist across aspects. Slightly higher accuracies were observed for categories that tend to be framed in direct evaluative terms, see examples (2) and (3).

(2) I wanna go help fight against Israel free Palestine.

(3) You have to pay for what you did with Palestine and you will beg to MUSLIMS to forgive them InshaAllah.

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 performance on the test set.

In example (2), the model correctly identifies the comment as expressing positive sentiment toward *Palestine*. When analyzed with respect to the aspect *Israel*, the same sentence is classified as negative, demonstrating the model's ability to assign distinct sentiment values to co-occurring targets. In contrast, aspects such as *Israel*, or *Zionists* often involve sentiment that is expressed indirectly through metaphor, irony, or euphemism. For example, the comments in (4) and (5) reflect negative sentiment through accusatory rhetorical contrast, a strategy that encodes stance implicitly.

(4) What Hitler did to Jews. Jews are doing the same to ISRAEL.

(5) One politician who is not afraid of the ZIONISTS.

An error analysis revealed three primary challenges: 1) the difficulty in detecting implicit sentiment, particularly in sarcastic or coded language; 2) the misclassification of sentences containing multiple possible targets, e.g., example (6); and 3) challenges in processing mixed-language comments and transliterations, where sentiment-bearing terms appeared in hybrid forms (e.g., Arabic-English transliterations). Our error analysis also shows the difficulty of capturing emotional nuance in politically charged and multilingual discourse, which is a major obstacle in analyzing user-generated content to real-world conflict events, especially when narrative framing relies on coded language or dog whistles.

(6) Free Jews from Israel

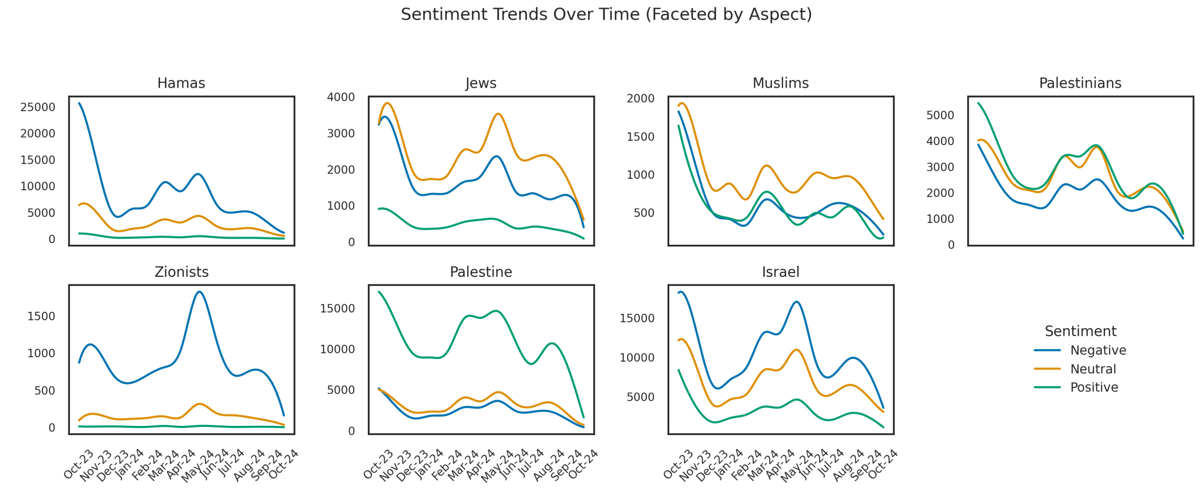

Figure 2: Sentiment trends over time across seven aspects. Each subplot shows smoothed counts of aspect-level sentiment predictions. Note: Y-axis scales differ across facets to account for large variation in volume of user-generated content. X-axis shows monthly aggregation.

5 RQ2: Using ABSA to Investigate Polarization

In order to determine whether we can use automatically annotated sentiment analysis to investigate polarization, we run the finetuned model on all segments that contain at least one aspect term in the complete dataset and analyze the results. While we do analyze polarization, our primary aim is to assess whether the proposed methodology can yield meaningful insights from large-scale datasets that are too extensive for manual analysis. We first show a longitudinal analysis of the sentiment per aspect and then an analysis of the sentiment across outlets.

5.1 Longitudinal Sentiment Dynamics

Figure 2 presents sentiment trends across seven aspects from October 7, 2023, to October 7, 2024. Emotionally charged discourse was most visible after the Hamas-led attacks, with sharp peaks in negative sentiment toward *Hamas*, *Israel*, and *Zionists*. Although overall comment volume declined after November 2023, engagement (and sentiment-carrying language) resurged in spring 2024, reflecting the temporal impact of external events.

Characteristics of these curves correspond to findings in earlier work that show increased social media activity around historical anniversaries (Royesh and Grossman, 2021; Miehling, 2024) and illustrate how sentiment differs in response to external trigger events. To show the relevance of the ABSA analysis for analyzing polarization in the comments, we investigate three sentiment peaks between October 2023 and May 2024 by aligning aspect-level sentiment with weekly comment volume and comparing them to the corresponding Shorts videos[10].

During the first peak (May 6-19), positive sentiment toward *Palestine* was primarily driven by user reactions to Shorts related to the Eurovision Song Contest, featuring repetitive expressions of solidarity such as “Free Palestine,” as well as to anti-Gaza war student protests, which included accusations of genocide against Israel. The second peak (May 13-26) coincides with Israel’s Independence Day. The peak reflects negative sentiment towards *Israel* shaped by discourse around identity and indigeneity, as reflected in trigrams such as “Canaan Palestinians descendants” and “indigenous people Palestinians”, together with accusations of genocide. The third peak, also in late May, targeted *Zionists* and was linked to campus protest reactions, with frequent terms such as “occupation colonization” and “Zionists declared terrorists”.

All of those findings of the longitudinal analysis show that we can find patterns that are known from prior work and that are interpretable given events during the time frame in which they occurred, thus corroborating our hypothesis that ABSA will be a useful tool for analyzing polarization in online discourse.

[10]For each peak, we looked at the most commented-on video IDs, reviewed their descriptions, and checked the corresponding video content to contextualize sentiment patterns.

5.2 Sentiment Across Outlets

Figure 3 highlights outlet-specific sentiment variation. Again, we can connect trends in this figure to prior work, thus corroborating that the ABSA analysis provides reliable insights into the polarization of the discourse. For example, *Palestine* is framed more positively in Middle Eastern outlets such as *TRT* and *Al Jazeera*. *BBC* and *DW* offer more mixed portrayals where most sentiment towards *Palestine* is positive or neutral. *Hamas* is framed negatively across all outlets, though sentiment strength varies; for *AJ*, *BBC*, and *TRT*, negative emotions range in the 71-73% range while for *DW*, the same emotion reaches 63%.

However, Figure 3 also shows that the aspect *muslims* is not represented in *Deutsche Welle (DW)* and appears a total of 12 times in the *Al Jazeera (AJ)* corpus (after dependency parsing). While the raw data contain 133 instances referencing *muslims*, the majority were excluded due to codeswitched content or syntactic filtering. This absence likely reflects a combination of platform-specific moderation policies, audience composition, and differing levels of salience in user-generated responses.

The aspect term *Zionists* consistently receives negative sentiment across all channels. However, this is based on a small sample size, since we only found a very limited number of mentions of this aspect. Thus, this analysis needs to be taken with a grain of salt. Here, a closer look at the comments is warranted.

Our results demonstrate that sentiment varies significantly across media outlets, confirming previously observed biases while offering new insights into polarization. User-generated discourse responding to content from Middle Eastern outlets such as *AJ* and *TRT* tends to frame *Palestine* and *Palestinians* more positively, contrasting with more mixed portrayals in *BBC* and *DW*. Negative sentiment toward *Hamas* is dominant across all platforms but varies in intensity. Our findings extend existing work that highlights the role of sentiment modeling in capturing ideological alignment and emotional intensity during geopolitical crises. For example, prior research has shown that shifts in affective tone in media coverage can predict or coincide with conflict events (Jamison et al., 2023), and that emotional exposure on social media can amplify polarization over time, particularly when communication occurs with out-groups and ideological alignment diverges (Lerman et al., 2024). Our aspect-level analysis refines this perspective by revealing how such dynamics manifest across ideological categories in the digital mainstream, capturing fine-grained sentiment distinctions within user-generated content; distinctions that are often obscured by global sentiment scores alone. Although our study analyzes polarization across multiple aspect terms, it also builds on the findings by Alamsyah et al. (2024), who examined the volume of pro-Israel and pro-Palestinian narratives on X (formerly Twitter) after October 7 and found the latter to be more prevalent. Our longitudinal sentiment analysis corroborates this imbalance and reveals how such patterns evolve across multiple outlets and intensify over the course of a year.

6 Conclusion and Future Work

We have investigated whether the use of aspect-based sentiment analysis of a large-scale corpus can provide useful insights for the study of polarization in discourse. We show that a current ABSA system can be successfully trained and used on more than 3.4 million YouTube comments we have collected. The intrinsic evaluation shows that the ABSA system reaches an F1 of 77.9, and the results are balanced between the three sentiment classes. Our extrinsic evaluation shows that the longitudinal analysis as well as the analysis across outlets corroborate known trends as well as provide more fine-grained or additional information that can be interpreted by a subject expert.

For the future, we are planning to evaluate the approach on a larger number of aspects. Since it is impossible for us to annotate enough data for training a model on all of those aspects, we will experiment with methods to determine the aspects automatically in combination with a zero-shot ABSA approach for the lower frequency aspects. Additionally, we are planning to use the large-scale sentiment annotations to investigate polarization in the Israel-Hamas war.

Limitations

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

Dataset Constraints: Several initial aspects (e.g., *Hezbollah*) were discarded because many

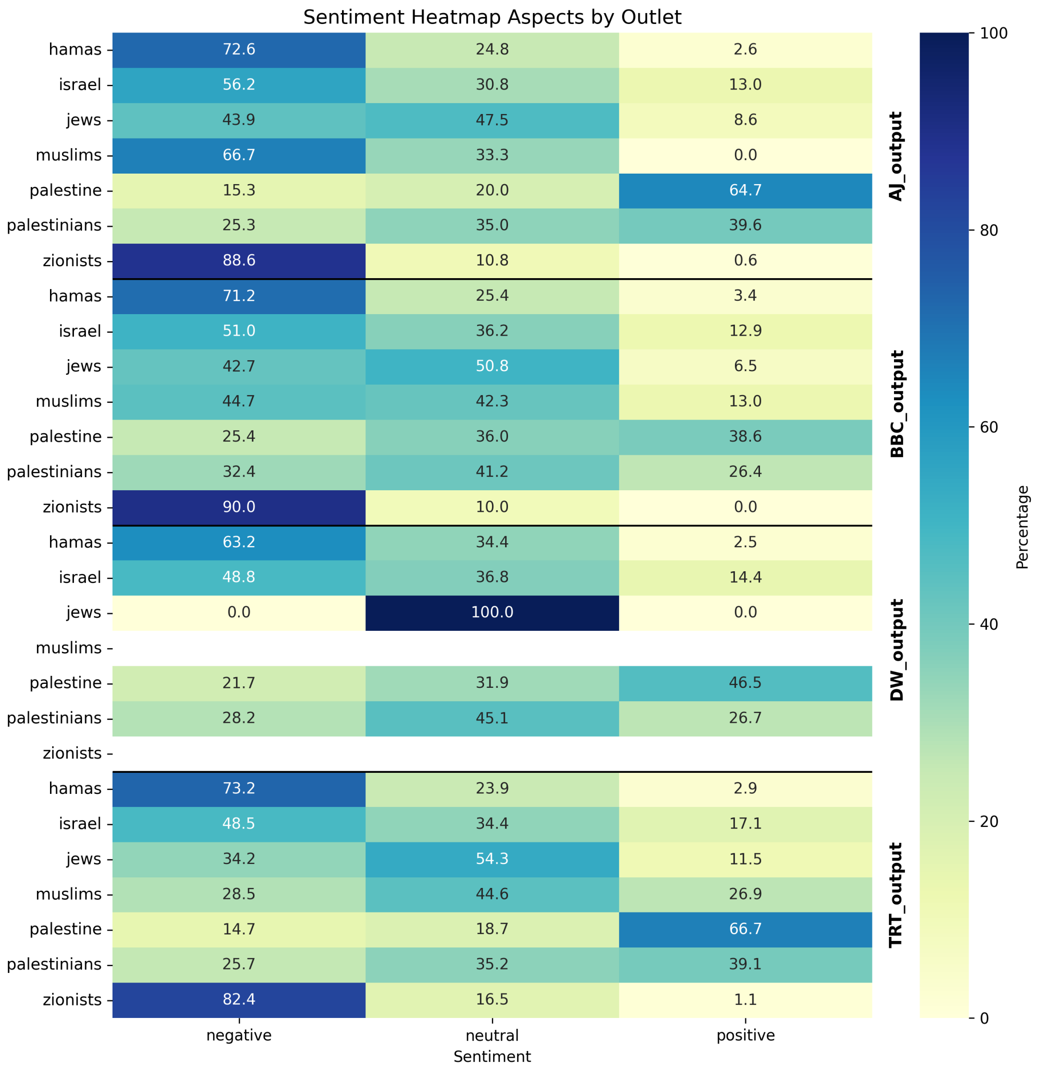

Figure 3: Heatmap for aspect sentiment across the four outlets.

comments failed the language threshold ($\geq 40\%$ English), indicating widespread codeswitching. Other aspects (e.g., *Zionists*) had very few positive examples, which led to class imbalance that impacted model performance. However, while data augmentation could mitigate some of these limitations, generating synthetic samples, particularly in sensitive political domains, raises ethical concerns.

Platform Coverage: The results on user-generated content from *Deutsche Welle (DW)* and *BBC* should be interpreted with caution: due to limited number of short-video content published by these outlets, our dataset contains significantly fewer samples from these sources compared to *TRT* and *Al Jazeera (AJ)*, reflecting outlet-specific content creation practices and not necessarily a lack of coverage of developments in our targeted topic.

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

Apoorv Agarwal, Boyi Xie, Ilia Vovsha, Owen Rambow, and Rebecca Passonneau. 2011. Sentiment analysis of twitter data. In *Proceedings of the Workshop on Languages in Social Media*, pages 30–38, Portland, Oregon. Association for Computational Linguistics.

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.

Miriam Amendola, Danilo Cavaliere, Cavaliere De Maio, Guiseppe Fenza, and Loia Vincenzo. 2024. Towards echo chamber assessment by employing aspect-based sentiment analysis and GDM consensus metrics. *Online Social Networks and Media*, 39-40:100276.

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.

Laleh Davoodi, Anssi Öörni, and Ville Harkke. 2024. Aspect based sentiment analysis of Finnish neighborhoods: Insights from Suomi24. In *Proceedings of the 9th International Workshop on Computational Linguistics for Uralic Languages*, pages 1–11, Helsinki, Finland. Association for Computational Linguistics.

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. Österreichische Akademie der Wissenschaften.

Tal Orian Harel, Jesscia K. Jameson, and Ifat Maoz. 2020. The normalization of hatred: Identity, affective polarization, and dehumanization on Facebook in the context of intractable political conflict. *Social Media + Society*, 6(2).

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *Preprint*, arXiv:2111.09543.

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.

Jordan Miner and John E. Ortega. 2024. An NLP case study on predicting the before and after of the Ukraine–Russia and Hamas–Israel conflicts. In *Proceedings of the Third Workshop on NLP for Positive Impact*, pages 140–151, Miami, Florida, USA. Association for Computational Linguistics.

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.

Pablo Muñoz, Alejandro Bellogín, Rodrigo Barba-Rojas, and Francisco Díez. 2024. Quantifying polarization in online political discourse. *EPJ Data Science*, 13(1):39.

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 fine-tuning for aspect-target sentiment classification. In

Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4933–4941, Marseille, France. European Language Resources Association.

Shamik Roy and Dan Goldwasser. 2021. Analysis of nuanced stances and sentiment towards entities of US politicians through the lens of moral foundation theory. In *Proceedings of the Ninth International Workshop on Natural Language Processing for Social Media*, pages 1–13, Online. Association for Computational Linguistics.

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.

Uddagiri Sirisha and Bolem Sai Chandana. 2022. Aspect based sentiment & emotion analysis with ROBERTa, LSTM. *International Journal of Advanced Computer Science and Applications*, 13(11).

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.

Cédric Treuillier, Sylvie Castagnos, Catherine Lagier, Ariane Grumbach, and Francesco Ricci. 2024. Gaining a better understanding of online polarization by approaching it as a dynamic process. *Scientific Reports*, 14(1):8702.

Vorakitphan Vorakit, Marco Guerini, Elena Cabrio, and Serena Villata. 2020. Regrexit or not Regrexit: Aspect-based sentiment analysis in polarized contexts. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 219–224, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Yu Wang, Tom Clark, Jeffrey Staton, and Eugene Agichtein. 2014. Towards tracking political sentiment through microblog data. In *Proceedings of the Joint Workshop on Social Dynamics and Personal Attributes in Social Media*, pages 88–93, Baltimore, Maryland. Association for Computational Linguistics.

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.

Wajdi Zaghouani, Mustafa Jarrar, Nizar Habash, Houda Bouamor, Imed Zitouni, Mona Diab, Samhaa El-Beltagy, and Muhammed AbuOdeh. 2024. The FIGNEWS shared task on news media narratives. In *Proceedings of the Second Arabic Natural Language Processing Conference*, pages 530–547, Bangkok, Thailand. Association for Computational Linguistics.

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.

Thomas Zeitzoff. 2016. Does social media influence conflict? Evidence from the 2012 Gaza conflict. *Journal of Conflict Resolution*, 62(1):29–63.

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*.