Masking the Bias: From Echo Chambers to Large Scale Aspect-Based Sentiment Analysis

Yeonjung Lee[1][0000-0003-4048-1841], Yusuf Mücahit Çetinkaya[1,2][0000-0001-5338-750X], Emre Külah[2][0000-0003-0877-5487], Hakkı Toroslu[2][0000-0002-4524-8232], and Hasan Davulcu[1][0000-0001-5602-8270]

[1] Arizona State University, Tempe, Arizona, USA
[2] Middle East Technical University, Ankara, Türkiye

Abstract. Aspect-based sentiment analysis (ABSA) is a natural language processing (NLP) task, ascribing precise sentiment linkages to specific entities and issues in text data. This paper addresses critical shortcomings in current ABSA methods, particularly the issues of limited aspects, training set biases, and lack of comprehensive stance-coded datasets. First, we develop a scalable MaskedABSA approach that masks aspect terms in training sentences to enable unbiased sentiment inference from the context alone. We show that the proposed method surpasses the state-of-the-art solutions in accuracy for the aspect term sentiment classification task, as verified by the SemEval datasets. Furthermore, we tackle the perennial challenges of limited training resources and the prohibitive costs of manual annotation in ABSA dataset creation by introducing an innovative weak supervision technique capitalizing on the inherent community clustering properties found within social media datasets. We utilize community detection algorithms to partition a share network into polarized groups with homogeneous adversarial stances, allowing large-scale aspect-based sentiment analysis dataset curation without labor intensive manual labeling. Our methodology is also validated using a real-world polarized dataset comprising diverse aspects and stances to showcase its efficacy and scalability.

Keywords: aspect-based sentiment analysis · social networks · weakly labeled data · social media.

1 Introduction

Aspect-based sentiment analysis (ABSA) is a subset of sentiment analysis that not only discerns the sentiment within a given text but also links it to specific entities or aspects. This precision enhances understanding of stance, enabling scholars to derive actionable insights by pinpointing the sentiments on specific entities or issues. As a result, ABSA has become a critical tool in data analytics and artificial intelligence (AI) applications.

While current ABSA methodologies have proven valuable, they require reevaluation to effectively handle the complexities and vast diversity present within large real-world datasets. One of the primary challenges of ABSA is training set bias, a phenomenon where the data used to train an algorithm fails to comprehensively represent the diversity of contexts and opinions in the full spectrum of the target populations. This bias can skew the algorithm's performance, often to the detriment of minority stances on less represented views, thereby compromising ABSA systems' efficacy and fairness.

In our model, we simply *mask* the terms corresponding to the aspects we seek to analyze. This technique allows our model to focus on the context alone to infer the sentiment directed towards the masked aspect without being biased by the aspect term itself. By doing so, we refine the model's capacity to more precisely gauge sentiments towards both previously seen and unseen (new) aspects, thereby addressing one of the critical limitations of current ABSA models.

The construction of labeled datasets for NLP tasks presents challenges due to the costly nature of labor-intensive manual coding. The high costs associated with ABSA dataset development results in a scarcity of publicly available datasets. ABSA differs from traditional sentiment analysis by marking both specific aspects present in every sentence and labeling the sentiment corresponding to each aspect. This process necessitates understanding the domain-dependent key aspects and the contextual interplay between aspects and intended sentiments, which cannot be easily automated. This requirement for manual labeling significantly contributes to the high costs associated with ABSA dataset development and the lack of publicly available datasets. The limited number and size of available datasets poses challenges for reliably training, testing and comparing models, especially transformer-based ones with millions of parameters. Researchers currently try to determine a small model's accuracy by averaging outcomes across multiple runs. However, larger datasets would facilitate more accurate evaluations [4].

Data analysis on social media platforms like X (formerly Twitter) is especially relevant due to such platforms' pervasive roles in enabling viral information cascades that impact public discourse. Recent studies [8],[3],[20] show that social media users pay attention to opinions they agree with. It indicates a strong correlation between biases in the content people both produce and consume. In other words, *echo chambers* reinforcing certain biased stances and a tailored media experience that eliminates opposing viewpoints and differing voices are very real on social media. Our approach employs community detection algorithms on retweet (or share) networks to reveal camps and their more homogeneous subcommunities. This method allows us to use partisan 'barrier-bound' users within a camp to rapidly and cost-effectively produce stance annotated weakly-labeled ABSA datasets.

For example, recent studies on political discourse on social media [3],[20] confirm that social media users are exposed mainly to opinions that agree with their own, and partisan users enjoy higher centrality and content endorsement. These findings indicate a strong correlation between biases in the content people both produce and consume. Our approach employs community detection algorithms on the retweet network, revealing camps and their more homogeneous sub-communities. Furthermore, by distinguishing between 'barrier-bound' partisan users, who only interact within their own camps vs. 'barrier-crossing' users [2,11], who attract positive engagement across camps, allows us to use 'barrier-bound' users and their shared core stances sub-community by sub-community to rapidly and cost-effectively produce weakly-labeled ABSA datasets. We summarize our key contributions as follows:

- We foster large-scale, cost-effective, weakly labeled ABSA datasets by harnessing homogeneous stances of partisan users in social media communities.
- We unveil MaskedABSA, a model that enhances aspect-based sentiment analysis by masking aspect terms in sentences. This improves accuracy in gauging sentiments toward both known and previously unseen aspects.
- We evaluate the MaskedABSA model's performance under imbalanced perspective conditions, particularly focusing on accuracies for detecting under-represented minority stances to bolster the fairness of ABSA systems.
- We make experimental datasets and models publicly available via GitHub at tweetpie/masked-absa.

The remainder of this article is structured as follows: Section 2 explores previous research relevant to our topic. Section 3 details the procedures involved in preparing the MaskedABSA and the weakly labeled dataset. Section 4 presents various experiments and evaluations. Section 5 offers concluding remarks and future works.

2 Related Works

ABSA has witnessed significant advancements, mainly through deep learning-based approaches and large-scale pre-trained models. E2E-ABSA [12] utilizes BERT for token classification with an additional layer for aspect polarity determination, yielding notable results. Dual-MRC [13] adopts a unique architecture featuring two adjacent BERT models, one each for extracting Aspect Terms (AT) and Opinion Terms (OT), which are combined to produce the outputs. Similarly, Span-ASTE [22] employs a span-level representation approach coupled with BERT to generate Aspect Terms and their corresponding stance pairs. Conversely, BART-Aspect-Based Sentiment Analysis Model (BART-ABSA) [23] is a sequence-to-sequence model developed to address ABSA subtasks, leveraging the BART model to tackle the Aspect-Category-Opinion-Sentiment (ACOS) subtasks. BART-ABSA [9] has demonstrated notable performance improvements through modifications in data pre-processing and model architecture, capitalizing on BART's pre-trained capabilities.

InstructABSA [18] diverges from the BERT architectures by employing pre-trained models and a sequence-to-sequence transformer by utilizing instruction queries. Our study uses the T5 transformer to train Masked Sentiment Classifier models. Introduced by C. Raffel et al. [10], the T5 Transformer demonstrates remarkable performance, particularly in question-answering tasks, compared to BERTBASE [6]. Our approach involved training the T5 model using a teacher forcing method, where the model is provided input sequences and their corresponding target sequences during training, resulting in enhanced performance in

ABSA tasks. Disentangled Linguistic Graph Model (DLGM) [15], and ABSA-DeBERTa [14], enrich BERT-based models' ability to capture intricate relationships between aspects and sentiments.

3 Methodology

3.1 Masked Aspect Sentiment Classification (MASC)

ABSA is a task that provides a polarity of each aspect term in the sentence. Therefore, Zhang et al. [24] define it as a comprehensive opinion summary at the aspect level. The process of aspect-based sentiment analysis comprises two distinct components: Aspect Term Extraction (ATE) and Aspect Term Sentiment Classification (ATSC). First, in the Aspect Term Extraction stage, the model identifies and extracts aspect terms from a provided sentence. Next, in the Aspect Sentiment classification phase, the model predicts the stance associated with identified aspect terms. As a result, the model returns aspect term and stance pairs. This task is illustrated in Figure 1. We formally define aspect-

Fig. 1: Example of aspect-based sentiment annotation.

based sentiment analysis process as follows. Consider a sentence S_i expressed as:

$$S_i = \{w_1, w_2, w_3, \ldots, w_n\} \tag{1}$$

where each w_j represents a word in the sentence. Within the sentence S_i, we identify multiple aspect terms denoted by A_i:

$$A_i = \{a_1, a_2, \ldots, a_k\},\ where\ k <= n \tag{2}$$

The process of detecting and extracting these aspect terms is termed Aspect Term Extraction (ATE). Subsequently, each aspect term is associated with a polarity P_i:

$$P_i = \{p_1, p_2, \ldots, p_k\},\ where\ k <= n \tag{3}$$

Each polarity p_x is categorized as one of the following values:

$$p_x = \{pos, neg, neutral, conflict, none\} \tag{4}$$

For the sample sentence in Figure 1, the model extracts aspects a_1 (price) and a_2 (service). It then determines the polarities as p_1 (positive) and p_2 (negative), respectively. As a result, the model outputs pairs such as (a_1, p_1) and (a_2, p_2), returning *(price, positive)* and *(service, negative)*.

During the Aspect Term Sentiment Classification task, we observe that current models tend to overemphasize the mere presence of the aspect term and the majority bias related to the term, neglecting the broader context of the sentence when determining sentiment. This tendency may not be readily apparent when utilizing SemEval datasets, as these are hand-crafted where training and test sets exhibit similar distributions. However, when deploying real-world datasets exhibiting highly skewed or missing labels for specific aspects, the vulnerability of current ABSA models becomes more evident. These models tend to memorize the associated word and repeatedly predict the sentiment label previously associated with it. To mitigate this issue, we introduce the MaskedABSA method. This technique involves replacing aspect terms within a sentence with a placeholder token, "[MASK]", thereby redirecting the model's focus towards the context of the sentence.

The procedure modifies the original sentence from $S_i = \{w_1, w_2, w_3, \ldots, w_n\}$ to $S'_i = \{w_1, a_1, w_3, \ldots, a_k, \ldots, w_n\}$ considering both the input S_i and aspect terms A_i. For the Mask method, we construct masked sentences such as $M_{i_1} = \{w_1, [MASK], w_3, \ldots, w_n\}$ and $M_{i_k} = \{w_1, w_2, w_3, \ldots, [MASK], \ldots, w_n\}$ from S'_i, where each aspect term a_k is replaced by "[MASK]" and for each aspect a new sample M_{i_k} is created. The model, denoted as F_{MASC}, is then tasked with predicting the polarity of the term represented by "[MASK]" in each sentence. The model returns a polarity value p_k. Thus, we can express the function of the model as $F_{MASC}(M_{i_k}) = p_k$.

3.2 Weakly Labeled Dataset

One effective strategy to address both the scarcity of ABSA datasets and the challenges of bias involves using weak supervision combined with social network structures, such as retweet networks [5]. Analyzing the patterns of retweets can infer the sentiment and aspect orientations prevalent within specific user groups or topics without requiring detailed sentence-by-sentence manual annotation. This approach allows for creating large-scale, weakly labeled datasets that can be refined using machine learning models and further validated through targeted reviews.

A well-defined codebook, offering guidelines for detecting and annotating stances on key aspects according to their distinct community orientations, significantly enhances the efficiency and consistency of the weakly labeled data production process.

Several researchers have utilized the Louvain community detection algorithm on social media networks [19],[1],[7], to extract non-overlapping communities. It was also observed that [25] users within the same Louvain community tend to share not only common interests but similar viewpoints and stances on popular issues and entities. Our approach harnesses per community homogeneous viewpoints, captured in a codebook, to generate vast amounts of weakly labeled datasets for ABSA tasks.

U.S. Race Relations Dataset Our dataset, derived from a real-world case, contains a collection of 9,084,824 tweets from January 1, 2022 to May 31, 2023, related to race relations in the U.S. The dataset involves 296,540 users and includes a retweet network with 738,012 edges, depicted in Figure 2. To collect the dataset, we used keywords related to U.S. race relations, specifically focusing on five topics - Black Lives Matter (BLM), All Lives Matter (ALM), Issues, Races, and Religion. Weakly labeled data is stance annotated data, leveraging shared community biases as recorded on the codebook'.

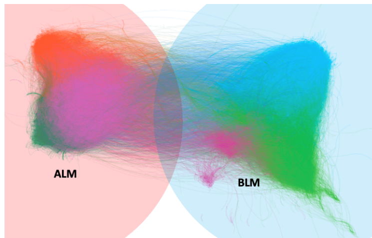

Fig. 2: Visualization of polarized camps in U.S. Race Relations

Fig. 3: Only Barrier-crossing users' network from U.S. Race Relations dataset

The primary advantages of producing weakly labeled datasets using our approach are its cost-efficiency and scalability, making it an appealing option in scenarios where acquiring comprehensive, fully labeled datasets is economically or logistically not feasible.

Camp Labeling of Users The retweet network graph shows a polarized network comprising two camps with seven sub-communities. Each sub-community is depicted in a different color in Figure 2.

An experimentally determined edge weight threshold is applied to partition the network into a pair of polarized camps. These camps are labeled by a panel of experts as either pro-BLM or pro-ALM, based on an analysis of their most viral 1,000 tweets, examined community by community. In social media analytics, users' communication patterns can be categorized as either barrier-bound or barrier-crossing. Barrier-bound partisan users interact only within their camps, whereas barrier-crossing users attract positive engagement from both camps. The barrier-crossing users' network shown in Figure 3. Previous studies [11],[17] have shown that contents amplified by barrier-bound users tend to reinforce existing biases within their groups, creating echo chambers that further polarize opinions. Our approach leverages the consensus that exists among barrier-bound users who exhibit clear partisan biases to develop a codebook and rapidly generate vast amounts of weakly labeled training datasets for MaskedABSA models.

Bias and Sentiment Coding We extracted all noun phrases and named entities from our corpus using weak labeling for significant aspect terms. Our codebook encompasses the most frequent and discriminating unigrams, bi-grams, and other n-grams, including all unambiguous bigrams such as "mental health" that could be overshadowed by more generic high-frequency unigrams like "health."

Subsequently, three domain experts independently coded the noun phrases and named entities for their shared stances within the BLM and ALM camps. These experts then convened in a panel to discuss and reconcile all their differences toward achieving a consensus on the codebook entries. From the top 3,000 phrases identified, 1,789 are determined to be non-ambiguous. The sentiment column reveals 496 phrases with negative connotations and 257 with positive ones. In terms of bias, 234 phrases are coded as "anti" and 160 as "pro" within the BLM camp; for the ALM camp, 177 phrases are "anti" and 90 are "pro".

The resulting dataset comprises weakly labeled tweets aimed at dissecting sentiment and ideological biases linked to specific camps and their subcommunities. Annotations comprise both aspect and stance as depicted in Figure 4.

Increased [funding/PRO] for [community programs/PRO] will [repair/PRO] the systemic injustices of [racial profiling/ANTI]

Fig. 4: An example of a stance annotated sentence from the BLM camp.

Preparing U.S. Race Relations and Politics Datasets From 62,196 samples, we separated two distinct sets of data focused on a pair of different themes: one on race-related issues and the other on politics-related issues and entities. Initially, we identified 1,206 terms associated with race-related topics and 593 terms with politics-related topics. We then created two separate collections of tweets: one matching terms related to race relations and another matching terms related to U.S. politics. We ensured these datasets did not overlap; any tweet that included terms from both categories was left out. After organizing these groups, we balanced the training and testing datasets for race relations and politics themes. We focused on ensuring fair representation from the All Lives Matter (ALM) and Black Lives Matter (BLM) camps. This approach led to the identification of 6,760 tweets for the politics dataset and 84,518 tweets for the race relations dataset.

4 Experiments and Evaluations

In our experiments, we tested our MaskedABSA approach using the SemEval Dataset and real-world datasets from Twitter.

4.1 SemEval Datasets

For the SemEval datasets, we used four training datasets: Laptop14, Restaurant14, Restaurant15, and Restaurant16. These datasets each include training and test subsets, allowing us to conduct our experimental evaluations comprehensively.

The data illustrated in Table 1 and Tabel 2 clearly demonstrate that the MaskedABSA model surpasses alternative models in terms of performance across all datasets, showing exceptional robustness even in cross-domain testing scenarios.

4.2 U.S. Race Relations and Politics Datasets

To ensure our dataset is similar in size and scope to the SemEval datasets, we randomly chose 3,000 tweets for training and 1,000 tweets for testing. While selecting these tweets for training and testing, we focused on maintaining a balance between the ALM and BLM camps. Due to the nature of Twitter, where negative messages often spread more effectively [21],[16], our dataset ended up having more negative content than positive. The details and key metrics of our dataset are as shown in Table 3.

Table 1: ATSC accuracy results (%) on SemEval datasets compared with MaskedABSA models.

Model	Laptop14	Res.14	Res.15	Res.16
BART-ABSA-ALSC[23]	76.76	*87.29*	76.49	89.10
Dual-MRC[13]	75.97	82.04	75.08	-
InstructABSA2 [18]	*81.56*	85.17	*84.50*	*89.43*
MaskedABSA	**86.24**	**87.65**	**91.53**	**94.83**

Table 2: Cross-domain ATSC results (%) on SemEval datasets comparing with the top model in the literature.

Train	Test	Model	Accuracy
Restaurant 14	Laptop14	InstructABSA	82.44
		MaskedABSA	**84.79**
Laptop 14	Restaurant 14	InstructABSA	80.56
		MaskedABSA	**85.32**
Restaurant 15	Hotel 15	InstructABSA	89.74
		MaskedABSA	**91.37**

Table 3: Race and Politics datasets distribution.

Races Dataset	Positive	Negative	Total	Politics Dataset	Positive	Negative	Total
Train	885	2115	3000	Train	873	2122	2995
Test	309	699	1000	Test	257	751	1008

We started by developing two models, one trained on the U.S. race relations dataset and the other on the U.S. politics dataset. The performance of each model was measured, with the race relations model achieving an F1 score of 71.92 and the politics model scoring 72.10, as shown in Table 4.

For cross-domain validation tests, we tested the performance of each model using the dataset from the other domain. The results showed only a small drops in their corresponding F1 scores, only 2.68 points for the U.S. politics dataset and 2.70 points for the U.S. race relations dataset. These experiments indicate that the MaskedABSA model maintains strong cross-domain performance, as shown in Table 4.

Table 4: Cross-domain test results on the U.S. Race Relations and Politics datasets using MaskedABSA.

Train	Single		Cross		Differences
	Test	F1	Test	F1	
Race	Race	71.92	Politics	69.24	-2.68
Politics	Politics	72.10	Race	69.40	-2.70

4.3 Bias Test

Aspect terms bias is a phenomenon where a model shows favoritism toward specific stances for specific terms. To tackle this problem, we developed the masking method, where context is forced to determine the stance instead of the aspect itself. To evaluate the efficacy of our approach, we conducted a bias test analyzing the accuracy of stance detection corresponding to messages carrying minority category biases. For instance, if the word "gun control" appears ten times in the dataset with eight occurrences labeled as anti/negative and two as pro/positive, we aimed to determine the accuracy of ABSA on the minority stance category tweets. In this test, we evaluated the performance of the InstructABSA model in comparison to the MaskedABSA model. The results demonstrated that our MaskedABSA model achieved an average accuracy of 45.6% and 57.7% in the politics and race-related models, significantly surpassing the InstructABSA's accuracies of 2% and 8% only, respectively. This experiment reveals a fundamental limitation of state-of-the-art ABSA models: a susceptibility to memorizing majority category stance labels for aspects, even in contexts that unambiguously signal a minority category stance. This highlights a critical weakness in their ability to accurately discern and represent the perspectives of less prevalent groups.

4.4 New Terms Test

To assess the stance detection accuracies for the previously unseen, new aspect terms issue, we designed an experiment to gauge the effectiveness of the MaskedABSA. We analyzed the model's performance with new aspect terms across U.S. race relations and politics datasets. Our approach involves replacing the queried aspect with a new term that does not appear in the training set. As a result, the aspect terms used in the test datasets were completely new to the models, having not been introduced during their training. This testing procedure was replicated for both the race relations and politics datasets to measure the performance of the InstructABSA model. Despite the experimental conditions, the MaskedABSA model's accuracy remained the same, as it conceals the queried aspect term during prediction, rendering the novelty of the term inconsequential. Additionally, we applied this methodology to the SemEval dataset, introducing new terms to challenge the model further. The results, documented in Table 5, clearly illustrate the MaskedABSA model's enhanced ability to handle new terms compared to the other models.

Table 5: Accuracy with replacements of new terms.

	Laptop 14	Res. 14	Res. 15	Res. 16	Race	Politics
Instruct (%)	50.00	71.23	55.08	64.40	53.00	61.84
Masked (%)	**86.24**	**87.65**	**91.53**	**94.83**	**71.92**	**72.10**

5 Conclusions and Future Work

In this study, we addressed the challenge of training set bias in ABSA models by introducing the MaskedABSA model, designed to focus on the contextual information of the text rather than merely memorizing majority term sentiments. Additionally, our study presents a methodology for constructing weakly labeled ABSA datasets by utilizing social network connectivity and a codebook comprising frequent terms. To ensure the robustness and generalizability of our approach, we rigorously evaluated it using SemEval benchmark datasets and real-world datasets on U.S. race relations and politics.

We conducted comprehensive tests to assess the model's effectiveness in handling previously unseen new aspect terms. Our findings demonstrated the superiority of the MaskedABSA model over the current state-of-the-art (SOTA) models. Rigorous cross-domain testing showed the model's versatility and robustness, showing only a marginal performance drop when applied to diverse real-world domains. The bias test confirmed our model's superior accuracy in mitigating aspect term bias, even for minority-held stances.

ABSA can be a powerful tool in detecting and countering mis/disinformation and fake news while navigating the complexities of breaking news. By dissecting the sentiment surrounding various aspects of a breaking news story, ABSA can identify inconsistencies in the language use to raise a red flag when the reported stances are inconsistent across credible mainstream vs. alternative media sources.

Furthermore, ABSA empowers users to explore the multifaceted landscape of breaking news by uncovering diverse perspectives from competing factions. By pinpointing communities engaged in discussions on social media and analyzing the sentiments within those groups, individuals can gain a deeper understanding of the varied viewpoints surrounding an event. This allows them to form well-rounded, informed opinions based on a broader spectrum of perspectives, transcending the limitations of their own echo chambers.

Acknowledgment

The work of Yusuf Mücahit Çetinkaya was supported by TUBITAK-2214-A.

References

1. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. Journal of statistical mechanics: theory and experiment **2008**(10), P10008 (2008)
2. Buhrmester, M.D., Cowan, M.A., Whitehouse, H.: What motivates barrier-crossing leadership? New England Journal of Public Policy **34**(2), 7 (2022)
3. Çetinkaya, Y.M., Toroslu, İ.H., Davulcu, H.: Developing a twitter bot that can join a discussion using state-of-the-art architectures. Social network analysis and mining **10**, 1–21 (2020)
4. Chebolu, S.U.S., Dernoncourt, F., Lipka, N., Solorio, T.: A review of datasets for aspect-based sentiment analysis. In: International Joint Conference on Natural Language Processing (2022)
5. Cherepnalkoski, D., Mozetič, I.: Retweet networks of the european parliament: Evaluation of the community structure. Applied network science **1**, 1–20 (2016)
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
7. Elmas, T., Randl, M., Attia, Y.: # teamfollowback: Detection & analysis of follow back accounts on social media. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 18, pp. 381–393 (2024)
8. Garimella, K., De Francisci Morales, G., Gionis, A., Mathioudakis, M.: Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In: Proceedings of the 2018 world wide web conference. pp. 913–922 (2018)
9. Hoang, C.D., Dinh, Q.V., Tran, N.H.: Aspect-category-opinion-sentiment extraction using generative transformer model. In: 2022 RIVF International Conference on Computing and Communication Technologies (RIVF). pp. 1–6 (2022)

10. Khashabi, D., Chaturvedi, S., Roth, M., Upadhyay, S., Roth, D.: Looking Beyond the Surface: A Challenge Set for Reading Comprehension over Multiple Sentences. In: Proc. of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL) (2018)
11. Lee, Y., Ozer, M., Corman, S.R., Davulcu, H.: Identifying behavioral factors leading to differential polarization effects of adversarial botnets. ACM SIGAPP Applied Computing Review **23**(2), 44–56 (2023)
12. Li, X., Bing, L., Zhang, W., Lam, W.: Exploiting BERT for end-to-end aspect-based sentiment analysis. In: Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)
13. Mao, Y., Shen, Y., Yu, C., Cai, L.: A joint training dual-mrc framework for aspect based sentiment analysis. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 13543–13551 (2021)
14. Marcacini, R.M., Silva, E.: Aspect-based sentiment analysis using bert with disentangled attention. LatinX in AI at International Conference on Machine Learning 2021 (2021)
15. Mei, X., Zhou, Y., Zhu, C., Wu, M., Li, M., Pan, S.: A disentangled linguistic graph model for explainable aspect-based sentiment analysis. Knowledge-Based Systems **260**, 110150 (2023)
16. Mousavi, M., Davulcu, H., Ahmadi, M., Axelrod, R., Davis, R., Atran, S.: Effective messaging on social media: What makes online content go viral? In: Proceedings of the ACM Web Conference 2022. pp. 2957–2966 (2022)
17. Rao, A., Morstatter, F., Lerman, K.: Retweets amplify the echo chamber effect. In: Proceedings of the International Conference on Advances in Social Networks Analysis and Mining. pp. 30–37 (2023)
18. Scaria, K., Gupta, H., Goyal, S., Sawant, S.A., Mishra, S., Baral, C.: Instructabsa: Instruction learning for aspect based sentiment analysis. arXiv preprint arXiv:2302.08624 (2023)
19. Seckin, O.C., Atalay, A., Otenen, E., Duygu, U., Varol, O.: Mechanisms driving online vaccine debate during the covid-19 pandemic. Social Media+ Society **10**(1), 20563051241229657 (2024)
20. Torres-Lugo, C., Yang, K.C., Menczer, F.: The manufacture of partisan echo chambers by follow train abuse on twitter. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 16, pp. 1017–1028 (2022)
21. Tsugawa, S., Ohsaki, H.: Negative messages spread rapidly and widely on social media. In: Proceedings of the 2015 ACM on conference on online social networks. pp. 151–160 (2015)
22. Xu, L., Chia, Y.K., Bing, L.: Learning span-level interactions for aspect sentiment triplet extraction. arXiv preprint arXiv:2107.12214 (2021)
23. Yan, H., Dai, J., Qiu, X., Zhang, Z., et al.: A unified generative framework for aspect-based sentiment analysis. arXiv preprint arXiv:2106.04300 (2021)
24. Zhang, W., Li, X., Deng, Y., Bing, L., Lam, W.: A survey on aspect-based sentiment analysis: Tasks, methods, and challenges. IEEE Transactions on Knowledge and Data Engineering (2022)
25. Zhang, Y., Wu, Y., Yang, Q.: Community discovery in twitter based on user interests. Journal of Computational Information Systems **8**(3), 991–1000 (2012)