BiasGanda at FIGNEWS 2024 Shared Task: A Quest to Uncover Biased Views in News Coverage

Al Manar Al Wardi,[1] **Blqees Al Busaidi,**[2] **Malath Al-Sibani**[1]
Hiba Salim Mohammed Al-Siyabi,[1] **Najma Al Zidjaly**[1]

[1]Department of English and Translation, SQU
[2]English Language, College of Education, SQU

{s138313,s138771,s137315,s125392}@student.squ.edu.om, najmaz@squ.edu.om

Abstract

In this study, we aimed to identify biased language in a dataset provided by the FIGNEWS24 committee on the Gaza-Israel war. We classified entries into seven categories: Unbiased, Biased against Palestine, Biased against Israel, Biased against Others, Biased against both Palestine and Israel, Unclear, and Not Applicable. Our team reviewed the literature to develop a codebook of terminologies and definitions. By coding each example, we sought to detect language tendencies used by media outlets when reporting on the same event. The primary finding was that most examples were classified as "Biased against Palestine," as all examined language data used one-sided terms to describe the October 7 event. The least used category was "Not Applicable," reserved for irrelevant examples or those lacking context. It is recommended to use neutral and balanced language when reporting volatile political news.

1 Introduction

This shared task aimed to detect bias and double standards in the news coverage of the Gaza-Israel war of 2023–2024, striving for a deeper understanding of these events. It seeks to build a cooperative community and educate the next generation of researchers in the field of Natural Language Processing (NLP) by establishing a comprehensively annotated shared corpus.

This paper adheres to the required title case format provided by the FIGNEWS24 team and aligns with the shared task's goal of framing the Israel-Gaza war of 2024. To achieve this objective, extant research was reviewed. For example, Recasens et al. (2013) conducted research to detect linguistic features of bias in a database derived from Wikipedia. Similarly, Gipp et al. (2021) developed annotation guidelines to assist annotators in understanding and detecting biased language. Additionally, Guo et al. (2022) introduced a framework utilizing masked-language models to detect biased language across different news outlets. They tested their framework on articles about five highly popular and diverse topics from ten different news outlets to analyze the language used to convey news.

This project does not only aim to uncover biased views in news coverage, but it also contributes to the ongoing research and development of NLP field methodologies. Specifically, Guo et al., (2022) employed quantitative methods, including a content-based approach that directly detected biased language from published sources by focusing primarily on the tone and word choices used when describing the same event across different news outlets. Hamborg (2020), in addition, emphasized the importance of word choice and labeling (WCL) in shaping audience perceptions of reported news. He defined WCL as: "When journalists refer to the same semantic concept by using different terms that frame the concept differently and consequently may lead to different assessments by readers, such as the terms 'freedom fighters' and 'terrorists,' or 'gun rights' and 'gun control'" (p. 79). Hamborg also explored various methods to reveal media bias to news consumers, aiming to mitigate the biased language impact in future reporting of similar events. He further suggested that most biased news reports originate from and are controlled by large media corporations within specific countries, citing the USA as an example: "In the US, for example, six corporations control 90% of the media" (79). Recasense et al., (2013) further provided

Proceedings of The Second Arabic Natural Language Processing Conference, pages 609–613

annotators with guidelines for the annotation process. In Framing the Israel-Gaza War in (FIGNEWS) shared task 2024, annotators were required to develop their guidelines based on general instructions set by the task's organizers.

Our team followed these instructions and developed guidelines informed by research in the field of natural language processing (NLP) and personal insights derived from the annotation database of the shared task.

2 Annotation Methodology and Examples

2.1 Development of Annotation Guidelines

To develop the guidelines, the team implemented a detailed process. Meetings were held at least once a week to ensure ongoing collaboration and refinement; the foundational elements of the guideline were established, with the FIGNEWS team providing the main points and texts.

Our team utilized two references to identify all categories. Categories were chosen to label biased texts if either framing bias or epistemological bias against a group or groups of people were detected in the text; these categories include "Biased against Palestine," "Biased against Israel," "Biased against both Palestine and Israel," and "Biased against others."

Conversely, a text was labelled "Unbiased" if neither of these biases was present. Texts with unspecified terms or inadequate context were classified as "Unclear." Texts not related to the conflict were categorized as "Not applicable."

Our annotators adhered to these categories, conducted quality checks, developed strategies for handling ambiguities, and regularly consulted with our team's supervisor to review their work. The guidelines were continually refined to best suit the project until all examples were thoroughly addressed.

2.2 Data Annotation Process

The annotation process commenced with the distribution of examples from the MAIN sheet among the four annotators, allocating 450 examples to each (B01 and B02), and from the IAA-1 and IAA-2 sheets, with 100 examples assigned to each annotator. Subsequently, efforts were made to standardize the application of the guidelines by selecting ten examples from the MAIN sheet to evaluate the implementation of the bias categories and the guidelines.

Following this standardization exercise, each annotator proceeded to annotate their designated examples independently, documenting their observations on language use across different languages. When encountering challenges, annotators sought input from their peers or the group's supervisor.

Additionally, annotators reported their findings on how certain languages described or referred to specific events or parties during the group's weekly meetings. This practice aimed to identify common observations among annotators to be incorporated into the project's overall findings.

2.3 Inter-Annotator Agreement (IAA) Analysis

To ensure consistency among annotators, the organizers of the Framing the Israel War on Gaza (FIGNEWS) shared task 2024 and implemented the Inter-Annotator Agreement (IAA) methodology. This approach required annotators to independently annotate specific texts without any discussion, thus, researchers could assess consistency and evaluate the quality of the guidelines.

The examples included in the IAA sheet consisted of notably lengthy texts, providing ample context to minimize the likelihood of divergent interpretations by annotators.

3 Team Composition and Training

The team consisted of four female Omani annotators, native Arabic speakers and second-language English learners, aged 18-24. One annotator is also learning French. All were undergraduate students at Sultan Qaboos University's Department of English and Translation, with specializations in Translation, Literature, and Education.

Supervision was provided by Dr. Najma Al Zidjaly, an Associate Professor in the Department of English and Translation at Sultan Qaboos University, who offered guidance throughout the project. For example, she conducted an online meeting to elucidate the definitions of propaganda and bias, highlighting the distinctions between the two terms, and provided instruction on annotation echniques. We engaged in discussions with her via a WhatsApp chat group, where she clarified some annotation process aspects, such as creating a codebook and approaching specific phrases and texts.

Likewise, the organizers of FIGNEWS24 offered support in the form of email announcements and correspondence as well as feedback for our draft. Moreover, annotators engaged in lengthy conversations to share insights and solve problems.

Workload distribution was equitable, and we held frequent meetings to strategize and review our progress. We wrote the first draft of our guidelines collaboratively and then proceeded to annotate the required batches of FIGNEWS24's database individually. Throughout the annotation process, guidelines were edited each time we had a meeting to discuss annotation problems and their solutions.

strategy. Thus, we decided to focus solely on completing the bias subtask.

Throughout the annotation process, we maintained a consistent approach and adhered to a shared guideline to ensure uniformity in annotations, despite the complex nature of the data.

Effective communication and cooperation among team members were key factors that contributed to the quality of our work. When encountering annotation-related challenges, such as ambiguities, we would refer to the guidelines/codebook and collaboratively develop solutions. The incorporation of examples and comparisons with related work significantly aided

Category / Language		Not applicable	Unclear	Unbiased	Biased against Palestine	Biased against Israel	Biased against both Palestine and Israel	Biased against others
English	NO	12	16	173	94	26	14	25
	%	3.33%	4.44%	48.06%	26.11%	7.22%	3.89%	6.94%
Arabic	NO	2	10	224	73	36	7	8
	%	0.56%	2.78%	62.22%	20.28%	10.00%	1.94%	2.22%
French	NO	7	34	133	93	57	12	24
	%	1.94%	9.44%	36.94%	25.83%	15.83%	3.33%	6.67%
Hebrew	NO	19	7	58	243	2	4	27
	%	5.28%	1.94%	16.11%	67.50%	0.56%	1.11%	7.50%
Hindi	NO	8	23	85	96	53	32	63
	%	2.22%	6.39%	23.61%	26.67%	14.72%	8.89%	17.50%
Total	NO	96	121	673	599	174	69	147
	%	2.66%	4.99%	37.38%	33.27%	9.66%	3.83%	40.83%

Table 1: Distribution of percentages across different languages and categories.

4 Task Participation and Results

Given that our team originally comprised six members, we initially assigned each three members a specific task (e.g., bias), while the remaining three were assigned to another task (e.g., propaganda). However, due to the departure of two teammates, our team was reduced to only four members, necessitating a revision of our original

our annotation process, thereby enhancing overall team performance.

5 Discussion

Two types of bias were identified in our analysis: framing bias and epistemological bias. Framing bias, as described by Gipp et al., (2021), occurs when a subject is presented from a specific perspective or "frame." An instance of this bias in

our dataset is the following text: "Palestinians walk among the rubble, as they inspect houses destroyed in Israeli strikes during the conflict, amid the temporary cease-fire between #Hamas and #Israel, at Khan Younis refugee camp in the southern #Gaza Strip on Monday." This text is biased against both Palestine and Israel because the term "conflict" is used to minimize the severity of the damage, whereas Palestinian and Israeli media often refer to the current situation as a "war."

Epistemological bias, as defined by Recasense et al. (2013), involves presenting a subject in a way that presupposes propositions to be true or false, subtly implied, asserted, or hedged in the text. In our dataset, several biased texts exhibit this type of bias, such as "WATCH: On the first day of the Israel-Hamas truce, Palestinian men attended Friday prayers inside a mosque wrecked by Israeli airstrikes." This text is biased against Israel because the word "wrecked" implies an outcome; an unbiased term would be "bombed" or "destroyed."

Category	No	Percentage
Not Applicable	48	2.67%
Unclear	90	5%
Unbiased	673	37.39%
Biased against others	147	8.17%
Biased against both Palestine and Israel	69	3.83%
Biased against Israel	174	8.17%
Biased against Palestine	599	33.28%
Total	1800	

Table 2: The seven bias categories, each with its percentage.

In some instances, bias is detected in a text that is neither against Palestine nor Israel, but against other parties. In such cases, the text is classified as "Biased against Others." For example, the text "Good evening, may God bless you, our brothers in Mesopotamia. A former Iraqi diplomat: A quarter of what Gaza has is what the Iraqis have. I spoke to Ghazi, and he told me that the number of hospitals in Gaza does not exist in a large country like Iraq! Following this shocking conversation in an interview with Makan Al Akhbar channel" is biased against the Iraqi government.

The second finding, as shown in Table 1, reveals notable patterns in the distribution of biased language categories across different languages. Apart from Hebrew, the category with the lowest percentage of occurrences is "Not Applicable," reserved for texts deemed irrelevant to the Israel-Gaza War 2023-2024. This discrepancy suggests a potential error in the data collection process.

In contrast, Hebrew sources exhibit the lowest percentage of texts categorized as "Biased against Israel." This outcome can be attributed to the predominance of Israeli media outlets among Hebrew sources, unlike other languages where media outlets are more diverse and multinational. For instance, Arabic news encompasses not only reports from Arabic-speaking countries but also those from foreign news agencies, including Israeli spokespersons.

Approximately more than a third of all texts (37.39%) are categorized as unbiased, with Arabic texts displaying the highest percentage of unbiased content. This phenomenon can be attributed to our methodology, which considered reported speech and quotes as biased or unbiased based on how the media outlet presented them, rather than the content of the quote itself. For example, a report Like “Imam Hassan Chalghoumi strongly criticizes the atrocities committed by ISIS, Hamas, including the killing of children and women and rape, stressing that there is no cure in Islam. We must say that this is forbidden and is not permissible according to Islamic law! Listen to him” is considered unbiased because the media outlet neutrally presents the speaker's views, despite the speaker's stance against Palestine.

As Table 2 shows, across all languages, texts categorized as "Biased against Palestine" outnumber those categorized as "Biased against Israel." This is noticed in Arabic texts as well. For instance, in Table 1, one-fifth of Arabic texts are biased against Palestine, while only a tenth are biased against Israel. This discrepancy can be attributed to the prevalent opposition to Hamas among most Arabic news outlets. According to our guidelines, no distinction was made between the two.

An analysis of our results from the Bias subtask evaluation reveals several key findings. To begin with, we annotated 2,200 data points in total, which shows our active participation, making a fine enough contribution to the Bias subtask. Despite not having the highest number of data points, our

work was noteworthy. Furthermore, achieving a Kappa score of 31.0 in the IAA Quality Track reflects our team annotators' consistency and adherence to the guidelines within our team. Moreover, we achieved a Macro F1 Average score of 29.5 in the Centrality Track that puts us in rank 4, which also reflects our team's consistency and adherence to the guidelines across all teams on B1 and B2.

Our final rank is the 6th, and it indicates that we performed exceptionally well compared to other teams in the FIGNEWS shard task.

6 Conclusion

In our paper, we described the development of our project and the guidelines for sub-task-1, drawing from the methodologies outlined in two key research papers by Gipp et al. (2021) and Recasense et al. (2013). We also discussed the methodology employed during annotation and presented the results of our annotations. The findings vary across different languages, yet the highest frequency category (e.g., "Biased against Palestine") remains consistent across all languages, underscoring the pervasive nature of media bias.

7 References

Gipp, B., Krieger, D., Plank, M., & Spinde, T. 2021. Towards A Reliable Ground-Truth For Biased Language Detection [Data set]. Zenodo. https://doi.org/10.5281/zenodo.4625151

Guo, X., Ma, W., & Vosoughi, S. 2022, May. Measuring media bias via masked language modeling. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16, pages 1404 -1408.

Hamborg, F. (2020, July). Media bias, the social sciences, and nlp: Automating frame analyses to identify bias by word choice and labeling. In *Proceedings of the 58th annual meeting of the association for Computational Linguistics: student research workshop*, pages 79-87.

Recasens, M., Danescu-Niculescu-Mizil, C., & Jurafsky, D. (2013, August). Linguistic models for analyzing and detecting biased language. In Proceedings of the 51st annual meeting of the Association for Computational Linguistics (volume 1: long papers), pages 1650-1659.

Wajdi Zaghouani, Mustafa Jarrar, Nizar Habash, Houda Bouamor, Imed Zitouni, Mona Diab, Samhaa R. El-Beltagy and Muhammed AbuOdeh . 2024. The FIGNEWS Shared Task on News Media Narratives, In Proceedings of the Second Arabic Natural Language Processing Conference (ArabicNLP 2024). Association for Computational Linguistics, Bangkok, Thailand.