

# Sina at FigNews 2024: Multilingual Datasets Annotated with Bias and Propaganda

**Lina Duaibes**  
Birzeit University  
1205358@student.birzeit.edu

**Areej Jaber**  
Palestine Technical  
University-Khadoorie  
a.jabir@ptuk.edu.ps

**Mustafa Jarrar**  
Birzeit University  
mjarrar@birzeit.edu

**Ahmad Qadi**  
7amleh Center  
ahmad@7amleh.org

**Mais Qandeel**  
ÖREBRO University  
Mais.Qandeel@oru.se

## Abstract

The proliferation of bias and propaganda on social media is an increasingly significant concern, leading to the development of techniques for automatic detection. This article presents a multilingual corpus of 12,000 Facebook posts fully annotated for bias and propaganda. The corpus was created as part of the FigNews 2024 Shared Task on News Media Narratives for framing the Israeli War on Gaza. It covers various events during the War from October 7, 2023 to January 31, 2024. The corpus comprises 12,000 posts in five languages (Arabic, Hebrew, English, French, and Hindi), with 2,400 posts for each language. The annotation process involved 10 graduate students specializing in Law. The Inter-Annotator Agreement (IAA) was used to evaluate the annotations of the corpus, with an average IAA of 80.8% for bias and 70.15% for propaganda annotations. Our team was ranked among the best-performing teams in both Bias and Propaganda subtasks. The corpus is open-source and available at <https://sina.birzeit.edu/fada>

## 1 Introduction

Since October 7, social media has been flooded with posts, articles, images, and videos related to the Israeli War on Gaza. Such posts are often divided by hate, bias, and fake news either in favor of or against one of the parties or by remaining neutral, see e.g., "Framing the Israeli War on Gaza" is a shared task on news media narratives (Zaghouni et al., 2024), which is part of the 2<sup>nd</sup> ArabicNLP conference. The task aims to create a multilingual corpus that unravels the layers of bias and propaganda within news articles in various languages.

Such shared tasks and datathons are crucial in the NLP community to foster collaboration and advance research in specific areas. Previous efforts, such as SemEval-2020 Task 11 (Martino et al.,

2020) and TSHP-17 (Rashkin et al., 2017) have provided valuable resources for propaganda detection in news articles. The dual focus of FigNews on bias and propaganda is a novel approach that addresses the evolving nature of misinformation on social media platforms. The detection of propaganda on social media is crucial (Darwish et al., 2021), as it can polarize public sentiment, foster violent extremism and hate speech, and eventually erode democracies and diminish trust in democratic procedures (Abuaiadah et al., 2017). Notably, only a few corpora have been recently built to address these issues. Recent work by (Hamad et al., 2023) involved establishing a Hebrew dataset comprising 15,881 tweets for detecting offensive language. This dataset was manually annotated with four labels: hate, abusive, violence, and pornographic. Their work focused on detecting hate speech in Hebrew tweets and implemented in *SinaTools* (Hammouda et al., 2024). Additionally, the *WojoodNER* Shared Task 2024 offered a new NER dataset related to the Israeli War on Gaza called *Wojood<sup>Gaza</sup>* (Jarrar et al., 2024). Other notable works include TSHP-17 (Rashkin et al., 2017), QProp (Barrón-Cedeno et al., 2019), and PTC (Da San Martino et al., 2019). TSHP-17 and QProp are document-level corpora, while PTC is a sentence-level corpus. While SemEval-2020 (Martino et al., 2020) Task 11 is similar to FigNews (Zaghouni et al., 2024) in its objective, they differ in their data sources and focus areas.

This paper describes our participation in the FigNews. Our contributions are:

- Annotated Corpus (12K FB posts) for bias and propaganda, in 5 languages.
- Annotation guidelines ensuring consistency and accuracy.

**Remark:** The corpus presented in this article

does not cover the genocide, ethnic cleansing, or starvation events as they mostly happened after collecting the corpus.

The article is organized as follows: Section 2 describes the methodology, 3 presents our team composition and training; Section 3 presents our participation and results; Section 5 analyzes some errors, and Section 6 concludes the paper.

## 2 Annotation Methodology

The objective of the task is to address the complex landscape of social media discourse related to the Israeli War on Gaza 2023-2024. The task organizers provided participants with 15k posts from verified Facebook accounts, selected between October 6, 2023, and January 31, 2024, using "Gaza" as a query keyword across 5 languages: Arabic, Hebrew, English, French, and Hindi. The dataset consists of 15 batches, each containing 1000 posts.

### 2.1 Annotation Guidelines

Our understanding of "bias" is based on the work done by the United Nations Committee on the Elimination of Racial Discrimination and the European Commission against Racism and Intolerance ([European External Action Service, n.d.](#)). We define the notations 'bias' and 'propaganda' based on the UN and EU accounts, as:

**Bias:** is generally understood as an inclination or prejudice towards or against a particular person or group, often in a way considered to be unfair. In other words, it is an unreasonable preference or dislike that prompts someone to behave in a discriminatory way, often based on unfair judgment. This bias is typically based on prohibited grounds of discrimination such as race, religion, language, nationality, ethnicity, social background, gender, and others.

**Classifications of Bias:** we adopted the same classes provided in the Shared Task: (1) Biased against Palestine, (2) Biased against Israel, (3) Biased against others, (4) Biased against both Israel and Palestine, (5) Not Applicable, (6) Unclear, and (7) Unbiased. We also introduced a new feature called "Type of Bias", which can be either: (a) *Explicit* (تحييز صريح) if it is obvious and evident in the post, (b) *Implicit* (تحييز ضمني) if it is clear but not evident in the post, and (c) *Vague* (تحييز مبهم) in case of indirect and ambiguous bias. This feature is important from a methodological viewpoint as it encourages the annotators to think more during

classification. If a post contains biased content but not in a direct way it can be accounted as implicit.

**Propaganda:** misleading ideas or statements that can distort the truth or omit facts to promote a specific political or social agenda. These ideas are typically published by media outlets. For example, propaganda can take the forms of exaggeration, minimization, spreading doubts, name-calling, labeling, or intentional vagueness. All these forms have the common intention to spread false information and obscure facts.

**Classifications of Propaganda:** We adopted the four classes provided in the Shared Task: (i) *Propaganda*, (ii) *Not propaganda*, (iii) *Not Applicable*, and (iv) *Unclear*.

Additionally, we added a new column to classify Propaganda into three types: (1) *Propaganda must be deleted*: if it contains evident harmful content that poses risks to the safety and security of individuals or groups; (2) *Propaganda may be deleted*: if we cannot easily judge whether it is propaganda, depending on a specific context; and (3) *Propaganda not to be deleted*: if it is not clear and lacks harmful consequences and therefore does not warrant deletion.

**Remark:** Since the data was collected from Facebook posts some cases contain quoted content (e.g. an unbiased post quoting biased content). It was established in the guidelines that a post should not be classified as bias or propaganda based on its quotation, but rather on the post itself.

**An Example** of the guidelines mentioned earlier regarding quoted content is as follows: " Hamas and Islamic Jihad spare no effort to exploit religious institutions for terrorist purposes," the IDF said in a statement. This post is annotated as unbiased because it is a direct quote and does not include any additional commentary or interpretation.

### 2.2 Inter-Annotator Agreement (IAA)

To evaluate the quality of our annotations, we used the *F1*-score and Cohen's Kappa ([Cohen, 1968](#)) to compute the agreement between the annotators. The results are shown in Table 1.

The task organizers allocated 100 posts (10%) from each batch for IAA, including 20 posts randomly selected from each language. Overall, we annotated 12,000 posts, resulting in an IAA dataset of 1,200 posts. These were distributed among our 10 annotators following this scheme: (1) each annotator received 240 posts, (2) each post was anno-

Annotators' pair		Biased				Propaganda			
		Cohen's kappa		F1_score_weighted		Cohen's kappa		F1_score_weighted	
		All	Binary	All	Binary	All	Binary	All	Binary
1	6	0.57	0.57	0.79	0.79	0.76	0.85	0.85	0.98
	4	0.76	0.77	0.53	0.56	0.33	0.25	0.58	0.93
	8	0.29	0.28	0.64	0.65	0.11	1	0.28	1
	2	0.64	0.64	0.82	0.85	0.72	0.62	0.8	0.91
3	10	0.8	0.78	0.86	0.9	0.67	1	0.77	1
	4	0.51	0.55	0.75	0.79	0.81	1	0.87	1
	6	0.89	0.93	0.96	0.98	0.37	0.78	0.59	0.97
	8	0.97	1	0.98	1	0.79	0.96	0.85	0.98
9	10	0.3	0.38	0.59	0.77	0.07	0.44	0.32	0.83
	2	0.79	0.81	0.87	0.92	0.79	0.86	0.785	0.93
	4	0.58	0.72	0.79	0.9	0.82	0.98	0.75	0.95
	6	-0.11	-0.09	0.54	0.59	0.18	0.57	0.48	0.85
7	8	0.94	0.97	0.98	0.98	0.93	1	0.95	1
	10	0.47	0.49	0.74	0.76	0.93	1	0.95	1
	2	1	1	1	1	0.85	0.93	0.91	0.97
	4	1	1	1	1	0.81	0.91	0.89	0.98
5	8	0.51	0.63	0.72	0.83	0.15	0	0.49	93
	10	0.87	0.85	0.95	0.95	0.92	1	0.95	1
	2	0.52	0.55	0.77	0.85	0.46	0.54	0.49	0.8
	6	0.39	0.45	0.65	0.75	0.05	0	0.34	0.91
Average		0.808	0.8515	0.623	0.6535	0.7015	0.9475	0.5725	0.733

Table 1: IAA for bias and propaganda annotations.

tated by two different annotators, and (3) the 240 posts assigned to each annotator were distributed among four other annotators. Consequently, each pair of annotators had 60 posts in common.

**All vs. Binary IAA:** to evaluate whether a (dis)agreement was dominated by a certain class, we mapped all labels into binary categories: (*Bias* or *NotBias* and *others*) and (*Propaganda* or *NotPropaganda* and *others*). Table 1 demonstrates no class dominance because All and Binary evaluations are close to each other.

Looking at all Cohen's scores in Table 1, the average is 0.808 for bias, which is a "very good" agreement, and 0.7015 for propaganda, which is a "good" agreement overall. Agreement on propaganda was more challenging but the results are enhanced when it is considered as a binary.

### 3 Team Composition and Training

**Team composition:** We assembled a team of 10 Master's students specializing in Law at Birzeit University, comprising 7 females and 3 males. All team members are native Arabic speakers with a good command of English.

**Training phase:** We began by selecting 200 posts to train all students in annotation. After training, each student was assigned 1,200 posts for annotation.

**Ensuring consistency** We held three workshops to ensure consistency to discuss guidelines, address

challenges, and resolve disparities. The first workshop involved an expert who reviewed the annotations and added comments for the annotators to address. In the second workshop, the annotators met with the expert to discuss his comments on the posts. In the final workshop, after reviewing their annotations compared to the expert's, they discussed the points of agreement and disagreement with him.

#### 3.1 Annotation process

**Annotation Phase:** The dataset consisted of 12 batches, comprising 10,800 posts from the Main sheet, and 1200 posts from the IAA sheet. The annotation was carried out in two phases:

- Phase One:** We distributed Batch01 and Batch02, each with 180 posts, among team members. To ensure consistency with the guidelines, an expert reviewed all student annotations for these batches and provided feedback.
- Phase Two:** we assigned each annotator 450 posts from two different batches. This step allowed us to complete the annotation of all 12 batches (i.e. 12k posts).

#### Set quality standards

To set quality standards among annotators, after the annotation process was complete, each pair of annotators who had annotated the same

Subtask	Track	1st Place	2nd Place	3rd Place
<b>Bias</b>	Guidelines	NLPColab	Eagles	Narrative Navigators
<b>Bias</b>	IAA Quality	NLPColab	JusticeLeague	<b>Sina</b>
<b>Bias</b>	Quantity	DRAGON	NLPColab	<b>Sina</b>
<b>Bias</b>	Consistency	The Lexicon Ladies	NLPColab	Narrative Navigators
<b>Propaganda</b>	Guidelines	NLPColab	Bias Bluff Busters	<b>Sina</b>
<b>Propaganda</b>	IAA Quality	NLPColab	<b>Sina</b>	The CyberEquity Lab
<b>Propaganda</b>	Quantity	NLPColab	<b>Sina</b>	The CyberEquity Lab
<b>Propaganda</b>	Consistency	NLPColab	Bias Bluff Busters	Sahara Pioneers/The CyberEquity Lab

Table 2: FIGNEWS 2024 shared task results.

data held meetings to review the selected posts they disagreed on. They discussed their differences, and if they reached an agreement, they would change the label accordingly. If they could not agree, they kept the original label.

## 4 Task Participation and Results

### 4.1 Results

Table 2 displays the final results provided by the shared task organizers. Our Sina team achieved the third and second place in the IAA Quality and Quantity tracks for the Bias and Propaganda subtasks, respectively. In addition to third place in Propaganda Guidelines.

Table 3 and Table 4 illustrate the distribution of the bias classes and types of bias across languages respectively. Table 3 shows that about 27% of the posts are biased against Palestine and 63% of the posts are unbiased. Most of the bias against Palestine originated from French posts. Table 4 gives more statistics about the types of bias. As shown in this table, most of the posts annotated as *Explicit* bias are in Hebrew.

For propaganda results, Table 5 illustrates the distributions of propaganda classes across languages, which shows that 31% of the posts (3333) are annotated as "Propaganda", and 66% (7084) are "Not Propaganda". The majority of the propaganda originated from French posts. Table 6 illustrates the distribution of the type of propaganda classes among languages. As shown in the table posts that were classified as propaganda must be deleted were in French with 348 posts.

## 5 Error Analysis and Discussion

Despite training and supervision, errors may arise from subjective interpretation, ambiguous guidelines, or complex content. We explored the errors

Class	Ar	En	He	Fr	Hi	Total
Biased Against Palestine	466	514	595	807	534	<b>2916</b>
Biased Against Israel	94	79	23	19	70	<b>285</b>
Biased against Both	6	7	11	6	14	<b>44</b>
Biased against others	42	28	53	39	49	<b>211</b>
Unbiased	1371	1486	1369	1212	1386	<b>6824</b>
Not applicable	49	7	17	20	25	<b>118</b>
Unclear	132	39	92	57	82	<b>402</b>
<b>Total</b>	<b>2160</b>	<b>2160</b>	<b>2160</b>	<b>2160</b>	<b>2160</b>	<b>10800</b>

Table 3: Distribution of bias classes across languages

Type of Bias	Ar	En	He	Fr	Hi	Total
<i>Explicit</i> (تحييز صريح)	394	336	<b>563</b>	412	388	<b>2093</b>
<i>Implicit</i> (تحييز ضمني)	199	217	265	236	269	<b>1186</b>
<i>Vague</i> (تحييز مبهم)	36	37	59	52	27	<b>211</b>

Table 4: Types of Bias

Class	Ar	En	He	Fr	Hi	Total
Propaganda	524	679	648	809	673	<b>3333</b>
Not propaganda	1484	1443	1447	1297	1413	<b>7084</b>
Not applicable	48	11	17	16	24	<b>116</b>
Unclear	104	27	48	38	50	<b>267</b>
<b>Total</b>	<b>2160</b>	<b>2160</b>	<b>2160</b>	<b>2160</b>	<b>2160</b>	<b>10800</b>

Table 5: Distribution of propaganda subtask classes across languages

Class	Ar	En	He	Fr	Hi	Total
Propaganda Must be deleted	192	191	266	<b>348</b>	277	<b>1274</b>
Propaganda May be deleted	524	488	382	461	396	<b>2059</b>
Propaganda not to be deleted	451	422	565	648	436	<b>2522</b>

Table 6: Types of Propaganda classes.

and noted:

1. False positives in bias annotations occurred when annotators marked neutral content as biased. For instance, the post: "Israel launched attacks on Syria on Nov 10 in response to a drone strike on Eilat. The IDF claimed it

attacked an organization responsible for the drone. Watch for more details." This news excerpt is informative and not biased.

2. Misclassification of propaganda: Some content was wrongly labeled as "must be deleted" propaganda despite lacking direct harmful implications. For example: "BREAKING: Israeli forces are causing massive destruction in Gaza, in response to a terrorist attack by Hamas. Image source: Middle East Eye post." While it is propaganda, it shouldn't be classified as "must be deleted."

## 6 Conclusion

This article presents our contribution to the FigNews 2024, where we annotated a multilingual corpus of 12,000 Facebook posts for bias and propaganda across five languages. We extended the annotation guidelines for better consistency and accuracy, providing a foundation for future work in detecting bias in social media. Our plans include expanding the corpus to cover more critical events of the war and leveraging neural and large Language models to automatically detect bias and propaganda on social media posts.

## Ethical Considerations

Given the sensitive nature of the topics and media narratives related to the Israel War on Gaza, our annotators, who are lawyers, have undergone extensive training to ensure careful and fair judgments. They meticulously review both Arabic and English translations to avoid any bias that might arise from machine translation.

## Limitations

We recognize the limitations in our annotation process. This is because of the subjective nature of identifying bias and propaganda in social media posts, and the sensitivity of the datasets involved.

## Acknowledgments

We would like to acknowledge the contributions of Nejira Softic during the formulation of the guidelines. We would like to also thank the Master's students specializing in Law and IT at Birzeit University for their help in annotating the datasets, especially Maram Shour, Belal Abu Zaina, Bayan Abu Alawi, Zainah Abughosh, Aya Al Dimasy, Doaa Abozena, Waad Alsheikh, Qassam Abu Hakmeh,

Aseel Mustafa, Basel Awwad, Omar To'Mallah, and Dyala Fakhouri, and as well as Prof. Reem Al-Botmeh for her support during the course. Also thanks to Palestine Technical University - Kadoorie for its support.

## References

- Diab Abuaiadah, Dileep Rajendran, and Mustafa Jarrar. 2017. [Clustering Arabic Tweets for Sentiment Analysis](#). In *Proceedings of the 2017 IEEE/ACS 14th International Conference on Computer Systems and Applications*, pages 499–506. IEEE Computer Society.
- Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Propopy: Organizing the news based on their propagandistic content. *Information Processing & Management*, 56(5):1849–1864.
- Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin*, 70(4):213.
- Giovanni Da San Martino, Yu Seunghak, Alberto Barrón-Cedeno, Rostislav Petrov, Preslav Nakov, et al. 2019. Fine-grained analysis of propaganda in news article. In *Proceedings of the 2019 EMNLP-IJCNLP Conference*, pages 5636–5646. ACL.
- Kareem Darwish, Nizar Habash, Mourad Abbas, Hend Al-Khalifa, Husein T. Al-Natsheh, Houda Bouamor, Karim Bouzoubaa, Violetta Cavalli-Sforza, Samhaa R. El-Beltagy, Wassim El-Hajj, Mustafa Jarrar, and Hamdy Mubarak. 2021. [A Panoramic survey of Natural Language Processing in the Arab Worlds](#). *Commun. ACM*, 64(4):72–81.
- European External Action Service. n.d. Human rights guidelines on freedom of expression online and offline. [https://www.eeas.europa.eu/sites/default/files/11\\_hr\\_guidelines\\_external\\_action\\_en.pdf](https://www.eeas.europa.eu/sites/default/files/11_hr_guidelines_external_action_en.pdf).
- Nagham Hamad, Mustafa Jarrar, Mohammed Khalilia, and Nadim Nashif. 2023. [Offensive Hebrew Corpus and Detection using BERT](#). In *Proceedings of the 20th ACS/IEEE International Conference on Computer Systems and Applications (AICCSA)*. IEEE.
- Tymaa Hammouda, Mustafa Jarrar, and Mohammed Khalilia. 2024. [SinaTools: Open Source Toolkit for Arabic Natural Language Understanding](#). In *Proceedings of the 2024 AI in Computational Linguistics (ACLING 2024)*, Procedia Computer Science, Dubai. ELSEVIER.
- Mustafa Jarrar, Nagham Hamad, Mohammed Khalilia, Bashar Talafha, AbdelRahim Elmadany, and Muhammad Abdul-Mageed. 2024. [WojoodNER 2024: The Second Arabic Named Entity Recognition Shared Task](#). In *Proceedings of the Second Arabic Natural*

*Language Processing Conference (ArabicNLP 2024)*, Bangkok, Thailand. Association for Computational Linguistics.

- G. Martino, Alberto Barrón-Cedeno, Henning Wachsmuth, Rostislav Petrov, and Preslav Nakov. 2020. Semeval-2020 task 11: Detection of propaganda techniques in news articles. *arXiv preprint arXiv:2009.02696*.

Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937. Association for Computational Linguistics.

Wajdi Zaghouni, Mustafa Jarrar, Nizar Habash, Houda Bouamor, Imed Zitouni, Mona Diab, Samhaa R. El-Beltagy, and Muhammed AbuOdeh, editors. 2024. *The FIGNEWS Shared Task on News Media Narratives*. Association for Computational Linguistics, Bangkok, Thailand.