

# Narrative Navigators at FIGNEWS 2024 Shared Task: New Frontiers in Bias and Propaganda Annotation Techniques

Maryam Al Emadi<sup>1</sup>, Jana El Messelmani<sup>2</sup>, Lyna Bermak<sup>3</sup>,  
Goumana Saleh Abdullah<sup>1</sup>, Esraa Sherqawi<sup>1</sup>, Anissa Jrad<sup>1</sup>,  
Zied Zouabi<sup>4</sup>, Wajdi Zaghrouani<sup>5</sup>

<sup>1</sup>Hamad Bin Khalifa University, Doha, Qatar

<sup>2</sup>AlJazeera Media Network, Doha, Qatar

<sup>3</sup>Carnegie Mellon University, Doha, Qatar

<sup>4</sup>Independent Researcher, Doha, Qatar

<sup>5</sup>Northwestern University in Qatar, Doha, Qatar

mryamalemedi@gmail.com , elmesselmanij@aljazeera.net , lbermak@andrew.cmu.edu , goab68111@hbku.edu.qa ,  
essh68661@hbku.edu.qa , ajrad@hbku.edu.qa , zo.zied@gmail.com , wajdi.zaghrouani@northwestern.edu

## Abstract

This paper presents our team's contribution to the FIGNEWS 2024 Shared Task, which involved annotating bias and propaganda in news coverage of the Israel-Palestine conflict. We developed comprehensive guidelines and employed a rigorous methodology to analyze 2,200 news posts from several official Facebook accounts of news websites in multiple languages. Our team, Narrative Navigators, achieved third place in both the Bias Guidelines and Bias Consistency tracks, demonstrating the effectiveness of our approach. We achieved an IAA Kappa score of 39.4 for bias annotation and 12.8 for propaganda detection. These findings and our performance underscore the need for enhanced media literacy and further research to counter the impact of biased and misleading information on public understanding of the conflict.

## 1 Introduction

In the context of the 2023-2024 war on Gaza, the concepts of "propaganda" and "bias" are crucial in media studies. Bias refers to a disproportionate leaning towards or against

certain ideas, entities, or nations, often leading to skewed representations (Auburn, 2023). These narratives frequently integrate propagandistic elements to sway attitudes and opinions (Jones, 2024; Kamalipour, 2010). Media outlets employ various techniques to mold contexts and foster ideologies that support their agendas (Paul & Elder, 2004; Davie, n.d.).

Propaganda is strategically used to manipulate public opinion through biased information to promote political objectives, significantly shaping domestic and international discourse (Jones, 2019). The need for Arabic-specific propaganda detection tools is emphasized by Firoj et al. (2022), who reveal the challenges posed by diverse linguistic and cultural contexts. Xiao and Alam (2023) address emotion manipulation and logical fallacies in Arabic digital content, highlighting the importance of balanced training models.

Deep learning approaches for analyzing propaganda techniques in news posts have been introduced by Martino et al. (2019), demonstrating effectiveness in recognizing propaganda. Combining machine learning with human judgments to improve bias detection accuracy is explored by Lazaridou (2020), showcasing the value of integrating diverse insights. Raza et al. (2024) detailed an annotation

and modeling process for bias detection, while Spinde et al. (2021) introduced the BABE methodology, using distant supervision by experts to train neural models in detecting biased news sentences.

The FIGNEWS 2024 Shared Task proposed by Zaghouani et al. (2024) aims to investigate media representation of the Israel-Palestine conflict by focusing on bias and propaganda in news posts across multiple languages. Our team, 'Narrative Navigators,' focused on the critical tasks of identifying biased content and detecting propaganda techniques in media articles. We developed and adhered to tailored annotation guidelines that allowed us to navigate through narratives with precision. This systematic application of our guidelines ensured consistent and reliable annotations.

Despite advancements in detecting biases and propaganda, challenges persist, especially in politically charged content. Potential biases in annotators' judgments due to demographic backgrounds and personal perceptions remain a concern. Zaghouani and colleagues have developed extensive annotation guidelines for various Arabic NLP tasks including punctuation (Zaghouani & Awad, 2016b), author profiling (Zaghouani & Charfi, 2018b), text analytics in social media (Zaghouani & Charfi, 2018c), diacritization (Zaghouani et al., 2016) and non-native Arabic texts (Zaghouani et al., 2015; Zaghouani et al., 2014). Their work on large-scale Arabic error annotation guidelines and framework (Zaghouani et al., 2014) is particularly relevant to our study, as it demonstrates the importance of well-defined annotation guidelines for ensuring data quality and consistency. While these studies focus on Arabic, we adapt their methodologies and insights to develop guidelines for annotating bias and propaganda in multilingual news posts related to the Israel-Palestine conflict in the FIGNEWS 2024 Shared Task.

This paper presents an overview of our participation, methodology, annotation process, and key findings. As Narrative Navigators, we encountered several challenges which we discuss in detail, proposing strategies to enhance the identification of bias and propaganda. Our efforts underscore the importance of advancing media literacy and fostering informed public discourse

on the conflict, true to our team's guiding principle of navigating narratives.

## 2 Methodology

The FIGNEWS 2024 Shared Task focuses on creating a comprehensive corpus to annotate media representations of the 2023-2024 Israel-Palestine conflict, aimed at uncovering biases and double standards in global and local news coverage. Our team, comprising seven members with diverse backgrounds, developed guidelines for annotating bias and propaganda. We analyzed news posts across different languages and cultural contexts.

### 2.1 Development of Annotation Guidelines

For Subtask 1, we focused on identifying biases by applying labels such as 'Biased against Palestine,' 'Biased against Israel,' and 'Unbiased.' We developed a set of key examples to guide the annotation process:

- Biased Against Palestine: Negative portrayals, such as suggesting irreversible actions by Israeli leaders to bias readers against Palestinians and omitting broader territorial claims in descriptions of militant activities.
- Biased Against Israel: Inaccurate depictions of Palestinian civilians as aggressors and misleading portrayals of property disputes to question the legitimacy of Israeli land claims.

These examples were crucial in refining our annotation process and ensuring specificity in identifying biases.

For Subtask 2, we adopted a detailed methodological framework informed by prior research and supplemented with its specific examples from related studies to ensure consistency and accuracy in annotations (El-Khaissi et al., 2022).

Given the sensitive nature of the conflict-related content, we maintained a high standard of ethical awareness throughout the annotation process. We conducted regular training sessions to reinforce the importance of neutrality and objectivity among team members. To facilitate deeper textual analysis, we employed AI tools, including ChatGPT, and held frequent team

discussions to refine our annotation protocols.

Text Segment	Annotator 1	Annotator 2
<p>“They are taking over the hospital.” A girl appears in a video purporting to be from #Al-Shifa_Hospital in Gaza and claims that Hamas controls it and prevents the displaced from leaving it. What is the truth about this video? And who is behind it?</p>	Unbiased	The tweet raises questions about the authenticity of the video and the motives behind it; it's difficult to determine if the tweet is biased.

Table 1: Example of Annotation Disagreement and Difficult case.

However, we found that ChatGPT's information was not always sufficient, as it was not fully updated regarding the ongoing war on Gaza. Cross-verification and AI-assisted analysis-maintained annotation quality. Workload distribution aimed to prevent fatigue, targeting a 95% Inter-Annotator Agreement (IAA) score for consistency.

## 2.2 Inter-Annotator Agreement (IAA) Analysis

The IAA analysis, conducted according to FIGNEWS guidelines, ensured annotation reliability using two sheets, IAA-1 and IAA-2, and daily reviews for consistency. Cohen's Kappa was employed to gauge inter-annotator agreement and highlight areas needing adjustment, maintaining objectivity and robustness in our annotations. This was crucial for developing accurate bias and propaganda detection models. Some discrepancies and complex cases, detailed in Table 1, underscore the subjective challenges in interpreting bias in text.

## 3 Team Composition and Training

### 3.1 Annotator Demographics

Our annotation team is composed of individuals with diverse backgrounds and educational experiences. All team members hold a university

degree, with some currently pursuing an MA in Digital Humanities and Societies. This educational foundation ensures they are well-equipped to handle detailed annotation tasks. The team's capabilities are further enhanced by the inclusion of a professional from Al Jazeera, who brings practical experience in media analysis. This diverse team composition allows for a comprehensive and nuanced approach to the annotation process.

### 3.2 Training Procedure

To ensure a consistent and accurate annotation process, we conducted thorough training sessions for all team members. At the project's outset, an online meeting was held to establish and refine our annotation guidelines through a collaborative review of a sample set of 10 news posts. This initial training session allowed team members to discuss and clarify any ambiguities or concerns regarding the guidelines and the annotation process.

Following the initial training, the team leader conducted one-on-one meetings with each team member to address specific issues in guideline adherence. These personalized training sessions were critical for ensuring that all team members consistently applied the annotation standards and had a clear understanding of the task at hand.

### 3.3 Team Coordination

Effective team coordination was essential for our annotation project's success. Team members stayed connected through continuous emails and regular online meetings, allowing for prompt resolution of any uncertainties and sharing of insights. We used a Google spreadsheet as a dynamic tool to record progress, feedback, and updates to our annotation guidelines. For instance, when "Emotive Propaganda" and "Biased Terminology" were identified, we quickly updated our guidelines and documented these changes in the spreadsheet, ensuring accurate and consistent annotations. Open communication, comprehensive training, and detailed documentation fostered a cohesive and efficient environment, enhancing the quality of our annotations.

## 4 Task Participation and Results

Our team actively participated in the [FIGNEWS 2024 Shared Task](#), which focused on the meticulous annotation of news posts related to the Israel-Palestine conflict across various languages. The task required a comprehensive, multilingual approach to identify instances of bias and propaganda in the news posts.

Our annotation strategies were informed by insights gleaned from academic literature and comparable tasks from previous years. By integrating these insights into our methodology, we were able to refine our annotation schemes and enhance our analytical techniques, ensuring that our methods were current, effective, and aligned with established best practices. This alignment with related work allowed us to robustly validate and compare our results, confirming the reliability and relevance of our findings.

### 4.1 Results and Performance

Our team, Narrative Navigators, achieved notable success in the FIGNEWS 2024 Shared Task, securing third place in two separate tracks across both Bias and Propaganda subtasks. This performance demonstrates the effectiveness of our approach and the quality of our annotations in multiple areas.

In the Bias Guidelines Track, Narrative Navigators ranked third with a Guidelines Score of 0.7497. This score is a composite of our Document Score, which received full marks (8 out of 8), and our IAA Kappa score of 39.4. Our strong performance in this track highlights the comprehensiveness and clarity of our annotation guidelines, as well as the consistency among our annotators.

For the Bias Consistency Track, we again secured third place. This track evaluated the centrality of our annotations compared to other teams, reflecting how well our judgments aligned with the overall consensus. Our performance here underscores the reliability and balance of our annotation approach.

In terms of specific metrics for the Bias subtask, Narrative Navigators achieved an IAA Kappa score of 39.4, an accuracy of 56.5%, and a Macro F1 Average of 45.5%. Our F1 Bias\* score, which indicates binary bias determination, was 70.5%. These scores were competitive,

particularly considering the subjective nature of bias annotation.

For the Propaganda subtask, while we didn't place in the top three, we still made significant contributions. Our team achieved an IAA Kappa score of 12.8, an accuracy of 54.5%, and a Macro F1 Average of 54.6%. The F1 Prop\* score, indicating propaganda detection, was 63.1%.

In terms of quantity, we contributed significantly to both subtasks, annotating 2,200 data points for Bias and another 2,200 for Propaganda, each across 2 Main batches and 4 IAA sets.

Compared to other teams, our performance was consistently strong, particularly in the Bias subtask. While top-performing teams like NLPColab achieved higher scores in several metrics (e.g., IAA Kappa of 78.8 for Bias), our results were comparable to other high-ranking teams. Our third-place finishes in two tracks put us among the top-performing groups in these challenging tasks.

Our participation and strong performance in both Bias and Propaganda subtasks demonstrate the versatility of our annotation approach and guidelines. The consistent top-three rankings in multiple tracks for Bias, coupled with our contributions to the Propaganda subtask, validate the quality and reliability of our work in these complex and nuanced tasks of annotating multilingual news posts.

## 5 Conclusion

Our participation in the FIGNEWS 2024 Shared Task (Zaghouani et al. (2024)), reveals significant findings: 42% of texts showed bias against Palestine, 15% against Israel, and 68% contained propaganda. These results align with studies like Jones (2019), who highlighted propaganda's role in manipulating public opinion, and Martino et al. (2019), who demonstrated effective deep learning approaches for detecting propaganda. The high prevalence of propaganda supports Firoj et al. (2022) and Xiao and Alam (2023), who emphasize the need for nuanced tools to address complex linguistic and cultural media content.

This study explored bias and propaganda in media coverage of the Israel-Palestine conflict using a comprehensive annotation approach. Results indicate a complex landscape: 29.2% of 129,800

data points showed bias against Palestine, 10.9% against Israel, and 46.1% contained propaganda. Such imbalanced coverage underscores the media's influence on public perception of the conflict.

Identifying bias and propaganda posed challenges, with some texts unclear or not applicable, highlighting the need for advanced detection methods. These findings stress the importance of media literacy and critical thinking among news consumers.

These findings underscore the necessity for balanced reporting and advanced NLP tools to enhance media literacy and public understanding of the Israel-Palestine conflict. Ongoing scrutiny of media practices, especially in conflict reporting, is crucial for ensuring fair, accurate, and unbiased coverage. Developing advanced NLP tools can enhance media literacy and public understanding of global issues, promoting responsible journalism and informed public discourse.

## Acknowledgments

The effort reported in this paper is funded by grant NPRP14C0916-210015 from the Qatar National Research Fund part of Qatar Research Development and Innovation Council (QRDI).

## References

- Chiraz El-Khaissi, Kartikeya Siddhant, AbdelRahim Elmadany, Sarah Jeblee, Maryam Al-Khouja, Motaz Saad, Muhammad Abdul-Mageed, and Kareem Darwish. 2022. Detecting Offensive Language on Arabic Twitter. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 105-118. Association for Computational Linguistics. <https://aclanthology.org/2022.wanlp-1.11.pdf>.
- Firoj Alam, Hamdy Mubarak, Wajdi Zaghouani, Giovanni Da San Martino, and Preslav Nakov. 2022. Overview of the WANLP 2022 Shared Task on Propaganda Detection in Arabic. arXiv preprint arXiv:2211.10057.
- Gavin Davie. Gatekeeping Theory. *Mass Communication Theory: from Theory to Practical Application*. [Online]. Available: <https://masscommtheory.com/theory-overviews/gatekeeping-theory/> [Accessed: 4<sup>th</sup> May 2024].
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. Fine-Grained Analysis of Propaganda in News Article. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5636–5646, Hong Kong, China. Association for Computational Linguistics. DOI: 10.18653/v1/D19-1565. Available: <https://aclanthology.org/D19-1565>.
- Konstantina Lazaridou, Alexander Löser, Maria Mestre, and Felix Naumann. 2020. Discovering Biased News Articles: Leveraging Multiple Human Annotations. In *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*, pages 1268–1277, Marseille, France, 11–16 May 2020. European Language Resources Association (ELRA). [CC-BY-NC]. Available: <https://aclanthology.org/2020.lrec-1.159.pdf>.
- Luke Auburn. 2023. Study of headlines shows media bias is growing. University of Rochester News Center. [Online]. Available: <https://www.rochester.edu/newscenter/study-of-headlines-shows-media-bias-growing-563502/> [Accessed: 3<sup>rd</sup> May 2024].
- Marc Owen Jones. 2019. The Gulf Information War| Propaganda, Fake News, and Fake Trends: The Weaponization of Twitter Bots in the Gulf Crisis. *International Journal of Communication*, vol. 13. [Online]. Available: <http://ijoc.org/index.php/ijoc/article/viewFile/XXX/XXXXXX>. [Accessed: 4<sup>th</sup> May 2024].
- Marc Owen Jones. 2024. Twitter post. April 12. "Full text of the tweet [6/ It's disturbing that @TheAtlantic would allow these two crucially important aspects of Israeli propaganda (and indeed, misinformation), to be published. By all means criticize Hamas, but not on the basis of non-credible or plausible narratives. I also find it bizarre that]." [Online]. Available: <https://twitter.com/marcowenjones/status/1778753389817409966>. [Accessed: 4<sup>th</sup> May 2024].
- Richard Paul and Linda Elder. 2004. How to Detect Media Bias & Propaganda In National and World News. (Original work published June 30, 2004). [Online]. Available: [https://www.criticalthinking.org/TGS\\_files/SAM-MediaBias.pdf](https://www.criticalthinking.org/TGS_files/SAM-MediaBias.pdf) [Accessed: 10<sup>th</sup> May 2024].
- Shaina Raza, Mizanur Rahman, and Shardul Ghughe. 2024. Dataset Annotation and Model Building for Identifying Biases in News Narratives. In *Text2Story 2024: Seventh International Workshop*

- on Narrative Extraction from Texts held in conjunction with the 46th European Conference on Information Retrieval, Glasgow, Scotland. [Online]. Available: <https://ceur-ws.org/Vol-3671/paper1.pdf>. [Accessed: 10<sup>th</sup> May 2024].
- Timo Spinde, Manuel Plank, Jan-David Krieger, Terry Ruas, Bela Gipp, and Akiko Aizawa. 2021. Neural Media Bias Detection Using Distant Supervision With BABE - Bias Annotations By Experts. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1166–1177, Punta Cana, Dominican Republic. Association for Computational Linguistics. DOI: 10.18653/v1/2021.findings-emnlp.101. Available: <https://aclanthology.org/2021.findings-emnlp.101>.
- Wajdi Zaghouni and Anis Charfi. 2018. Annotation Guidelines for Text Analytics in Social Media. In *Qatar Foundation Annual Research Conference Proceedings*, Vol. 2018, Issue 3, Mar 2018, ICTPD879. Hamad bin Khalifa University Press (HBKU Press). DOI: <https://doi.org/10.5339/qfarc.2018.ICTPD879>.
- Wajdi Zaghouni and Anis Charfi. 2018. Guidelines and Annotation Framework for Arabic Author Profiling. In *Proceedings of the 5th Workshop on Open-Source Arabic Corpora and Processing Tools*. [http://rec-conf.org/workshops/lrec2018/W30/pdf/5\\_W30.pdf](http://rec-conf.org/workshops/lrec2018/W30/pdf/5_W30.pdf)
- Wajdi Zaghouni and Dina Awad. 2016. Toward an Arabic Punctuated Corpus: Annotation Guidelines and Evaluation. In *Proceedings of the 2nd Workshop on Arabic Corpora and Processing Tools 2016 Theme: Social Media*, page 22.
- Wajdi Zaghouni, Behrang Mohit, Nizar Habash, Ossama Obeid, Nadi Tomeh, Alla Rozovskaya, Noura Farra, Sarah Alkuhlani, and Kemal Oflazer. 2014. Large Scale Arabic Error Annotation: Guidelines and Framework. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages not listed, Reykjavik, Iceland. European Language Resources Association (ELRA). Available: [http://www.lrec-conf.org/proceedings/lrec2014/pdf/956\\_Paper.pdf](http://www.lrec-conf.org/proceedings/lrec2014/pdf/956_Paper.pdf).
- Wajdi Zaghouni, Houda Bouamor, Abdelati Hawwari, Mona Diab, Ossama Obeid, Mahmoud Ghoneim, Sawsan Alqahtani, and Kemal Oflazer. 2016. Guidelines and Framework for a Large Scale Arabic Diacritized Corpus. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 3637–3643.
- Wajdi Zaghouni, 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.
- Wajdi Zaghouni, Nizar Habash, Behrang Mohit, Abeer Heider, Alla Rozovskaya, and Kemal Oflazer. 2014. Annotation Guidelines For Non-native Arabic Text In The Qatar Arabic Language Bank. In *Qatar Foundation Annual Research Conference Proceedings*, Vol. 2014, Issue 1, Nov 2014, ITPP1105. Hamad bin Khalifa University Press (HBKU Press). DOI: 10.5339/qfarc.2014.ITPP1105.
- Wajdi Zaghouni, Nizar Habash, Houda Bouamor, Alla Rozovskaya, Behrang Mohit, Abeer Heider, and Kemal Oflazer. 2015. Correction Annotation for Non-Native Arabic Texts: Guidelines and Corpus. In *Proceedings of the 9th Linguistic Annotation Workshop*, pages 129–139, Denver, Colorado, USA. Association for Computational Linguistics. DOI: 10.3115/v1/W15-1614. Available: <https://aclanthology.org/W15-1614>.
- Yahya R. Kamalipour. 2010. Language, media and war: Manipulating public perceptions. *Media*, 3(1-4):87.
- Yang Xiao and Firoj Alam. 2023. Nexus at ArAIEval Shared Task: Fine-tuning Arabic language models for propaganda and disinformation detection. arXiv preprint arXiv:2311.03184. Available: <https://arxiv.org/abs/2311.03184>.

## Appendices

### Appendix A: Examples of Text Annotations and Inter-Annotator Agreement

Text Segment: “Hey Enoch. My partner in the military, Capt. Or Yosef Ran, was killed on Saturday when he arrived as the first force in Kfar Gaza. Orosh (that's what everyone called him) was one of the founders of the migration unit in Dovdvan, he had a T-shirt and was a heroic soldier, the son of Avri Ran, one of the founders of the settlements in Judea and Samaria. I would appreciate it if you could tell me a little about him. May everyone know this amazing soul, a free and kind soul, a good commander for his soldiers and always surrounded by many friends and family, bringing light wherever he went. This man a month ago told me that "if I die then only for this country" and on Saturday he left my house and never came back. Today we

will stand at "Givaot Olam", the hill where he grew up and loved so much, at 3:00 p.m. and we will pay our last respects to him." Annotator1: Not Propaganda Annotator 2: Propaganda.

**Notes:** Annotators disagreed on whether the text segment is propaganda, reflecting differing interpretations of its intent and emotional tone.

### **Appendix B: Example of Difficult Annotation Case**

Text Segment: "IDF fighter and MDA volunteer Lt. Col. Elkana Neulander, the 14th, fell in the battles in the Gaza Strip: "a man full of joy in life who will be greatly missed." The Magen David Adom family bows their heads and mourns the fall of an MDA volunteer medic, an IDF fighter. To Lt. Col. Elkana Newlander, who fell in the Gaza Strip, 24 years old in the fall. The late Elkana was a dedicated and professional medic, who started his career at the MDA in 2015, and since then he has been volunteering at the Jerusalem MDA station together with his father, MDA medic Avery Neulander. During his years at the MDA, he treated and saved the lives of many. Moshe Benita, a paramedic at the Jerusalem MDA station and Elkana's friend: "He was a great volunteer, it was always important for him to know everything that was happening in order to contribute and save lives at every opportunity. Adam's life was against his eyes and I always wanted to work with him. Elkanah was a man of character who behaved with sensitivity and devotion to all those around him. We will miss his smile and his big heart." Eli Gartner, head of volunteers at the Jerusalem MDA station: "The late Elkana was a man full of joy in life, a dedicated medic, very active, professional and full of compassion. He was an integral part of the Jerusalem MDA station. He was an amazing friend, there is no way to explain this, we will miss him very much." MDA CEO Eli Bin Spad to the late Elkana: "The late Elkana was a professional and dedicated MDA volunteer, who gave personal attention to each of his patients. Magen David Adom bows his head and hugs his father MDA volunteer Avery Neulander, his mother Nava, and his sisters and brothers, may his memory be blessed."

Annotator 1: Biased Against Palestine and Propaganda

Annotator 2: Unbiased, Propaganda.

**Notes:** This text segment presents a difficult case for annotation due to its emotionally charged content and potential bias. While one annotator viewed it as biased against Palestine and propaganda, the other was uncertain about the bias but agreed it was propaganda. This example illustrates the challenges in determining bias and intent in complex narratives.