

# Navigating Dialectal Bias and Ethical Complexities in Levantine Arabic Hate Speech Detection

Ahmed Haj Ahmed<sup>1</sup>, Rui-Jie Yew<sup>2</sup>, Xerxes Minocher<sup>1</sup>, Suresh Venkatasubramanian<sup>2</sup>,

<sup>1</sup>Haverford College, <sup>2</sup>Brown University

Correspondence: ahajahmed@haverford.edu

## Abstract

Social media platforms have become central to global communication, yet they also facilitate the spread of hate speech. For underrepresented dialects like Levantine Arabic, detecting hate speech presents unique cultural, ethical, and linguistic challenges. This paper explores the complex sociopolitical and linguistic landscape of Levantine Arabic and critically examines the limitations of current datasets used in hate speech detection. We highlight the scarcity of publicly available, diverse datasets and analyze the consequences of dialectal bias within existing resources. By emphasizing the need for culturally and contextually informed natural language processing (NLP) tools, we advocate for a more nuanced and inclusive approach to hate speech detection in the Arab world.

**Warning:** The content of this paper may be upsetting or triggering to some readers.

## 1 Introduction

In the Levant, deep-rooted socio-political tensions have turned language into a weapon. The rise of digital platforms has amplified hate speech, necessitating robust detection and mitigation mechanisms (Castaño-Pulgarín et al., 2021; Awan, 2014). Automated tools leveraging NLP are essential for curbing online hate speech (Jahan and Oussalah, 2023). However, these tools are not equally effective across all languages and dialects. While significant progress has been made for languages like English, Levantine Arabic remains underserved (Bender, 2019).

Levantine Arabic, spoken across Syria, Jordan, Palestine, and Lebanon, is a dialect continuum with significant regional variations, making it challenging for current NLP technologies to capture (Haff et al., 2022). Existing hate speech detection models often overlook the rich cultural and sociolinguistic nuances of the dialect. This paper addresses the

ethical, cultural, and linguistic challenges in detecting hate speech in Levantine Arabic and highlights the critical need for more representative datasets.

## 2 The Linguistic Complexity of Levantine Arabic

### 2.1 Dialectal Variation

Levantine Arabic is a continuum of dialects differing significantly across countries and regions. In Syria, the Damascus dialect contrasts with that of Idlib or rural areas; for instance, "clothes" is "awaei" in Damascus but "teyab" in Aleppo, and "girl" is "bint" in Damascus and "sabiye" elsewhere, with pronunciation variations like consonant softening altering meanings (Naïm, 2012). Jordanian Arabic varies between urban centers like Amman and rural areas that preserve traditional forms; "now" is "halla" in urban settings and "hassa" in rural regions, and the letter jim may be pronounced as a soft "j" or a hard "g" (Sakarna, 2005). Palestinian Arabic differs between Jerusalem, the West Bank, Gaza, and diaspora communities; "cup" is "kaseh" in Jerusalem but "kubayeh" in Gaza. In Lebanon, Beirut's Arabic incorporates French loanwords due to historical influences—unlike regions like Tripoli or the south; the pronunciation of the letter qaf also varies between a glottal stop, a hard "k," or the standard "q" sound (Obégi, 1971; Naïm, 2012).

These regional differences are deeply tied to cultural and socio-political identities. Variations in expressions, idioms, and pronunciation can carry different meanings depending on locality, posing significant challenges for NLP tools. Generic models, often trained on standardized Arabic, may not capture these subtleties

### 2.2 The Role of Sociolinguistic Context

Understanding hate speech in Levantine Arabic requires not only linguistic proficiency but also a deep understanding of the socio-political context

in which the language is used. The Levant is a region marked by ongoing conflicts, occupation, and political instability. Hate speech is often employed strategically to exacerbate sectarian divisions, mobilize political supporters, or criticize opposition groups.

In Syria, for instance, even subtle linguistic features like the pronunciation of the qaf have become sociopolitical markers. Historically a neutral phonetic variation, the qaf pronunciation shifted during the conflict to signal regime alignment (Omran, 2021). Security forces used it in propaganda to stoke sectarian fears, while opposition groups mocked it as a regime identifier, transforming a simple linguistic trait into a symbol of allegiance and deepening societal divides.

Similar dynamics can be observed in Lebanon, where political factions often use divisive rhetoric to maintain control. Hate speech is not merely offensive language but part of broader strategies to sustain political dominance and suppress dissent. Any attempt to detect and mitigate hate speech in this context must account for these complex and shifting dynamics, including the sociopolitical significance of linguistic nuances.

### 3 The Problem with Current Datasets

#### 3.1 Lack of Publicly Available Datasets

One of the most significant barriers to improving hate speech detection in Levantine Arabic is the lack of publicly available datasets. While several datasets exist for Modern Standard Arabic (MSA), Egyptian Arabic, Gulf Arabic, and others (Alakrot et al., 2018; Mubarak et al., 2017; Albadi et al., 2020; Al-Ajlan and Ykhlef, 2018), there is a striking absence of resources dedicated to Levantine Arabic. This gap limits the ability of researchers and developers to create effective hate speech detection models for the region.

The few datasets that do exist for Levantine Arabic are often restricted in scope, limiting their utility for broader research. Moreover, these datasets are rarely representative of the full spectrum of dialectal variation found within the Levant. Without publicly available, diverse datasets, the development of inclusive and effective NLP tools remains out of reach (Barocas et al., 2023).

#### 3.2 Dialectal Bias in Existing Datasets

Even the best available datasets for Levantine Arabic are biased toward specific regional dialects.

A prominent case in point is the Levantine Hate Speech and Abusive Language (L-HSAB) Twitter dataset—the first and only publicly available dataset dedicated to hate speech and abusive language in Levantine Arabic (Mulki et al., 2019). While L-HSAB is invaluable due to its size and scope, it disproportionately focuses on Lebanese Arabic. This bias stems primarily from its data collection methodology, which involved extracting tweets using keywords related to Lebanese political figures and events (Barocas and Selbst, 2016).

The most frequently mentioned entities in L-HSAB are predominantly Lebanese. "Gebran Bassil," a Lebanese politician, is mentioned over 1,000 times. The term "Lebanon" appears around 250 times, and "Wiam Wahhab," another Lebanese politician and journalist, is mentioned approximately 200 times. This concentration on specific individuals and topics skews the dataset toward Lebanese political discourse, thereby overlooking the linguistic and sociopolitical nuances present in other Levantine regions.

This skew introduces significant bias, as the linguistic features, idiomatic expressions, and even manifestations of hate speech in Lebanese Arabic differ markedly from other Levantine dialects. For instance, certain derogatory terms or politically charged phrases common in Lebanese discourse may be absent or hold different connotations in Syrian or Jordanian contexts. A term like "za'ran", meaning "thugs" in Lebanese Arabic, is a strong insult in Lebanon but does not carry the same weight in Syrian Arabic. Conversely, a Syrian expression such as "shabbiha", referring to pro-regime militias, is a loaded term in Syria but might not evoke the same response or understanding among Lebanese speakers (Ümit Üngör, 2020).

Moreover, the focus on specific events and actors further narrows the dataset's applicability. The political landscape and issues prevalent in Lebanon are unique and may not reflect the concerns or conflicts in Syria, Jordan, or Palestine. Hate speech related to Lebanese political parties like the Free Patriotic Movement or events like the Lebanese protests of 2019 would not encompass the types of hate speech prevalent in other regions.

As a result, models trained on datasets like L-HSAB are less likely to generalize effectively to other dialects. They may fail to detect hate speech in Syrian, Jordanian, or Palestinian Arabic due to differences in vocabulary, idioms, and sociopolitical references. This limitation reduces the overall

effectiveness of hate speech detection tools across the Levantine region.

Furthermore, this bias can lead to misclassification, where non-hateful speech in one dialect is incorrectly flagged as abusive because the model does not accurately interpret the linguistic nuances of that dialect. Conversely, actual hate speech may go undetected in underrepresented dialects, allowing harmful content to proliferate.

In summary, while datasets like L-HSAB are crucial stepping stones in advancing hate speech detection for Levantine Arabic, their dialectal and topical biases highlight the need for more inclusive data collection strategies. Expanding the dataset to include a broader range of dialects and sociopolitical contexts is essential. By doing so, we can develop NLP tools that are both effective and equitable, ensuring that all communities within the Levantine region are adequately represented and protected in the digital space (Barocas et al., 2023).

### 3.3 Limitations of Pre-trained Embeddings and the Need for Domain-Specific Models

In addition to dataset biases, the choice of language models and embeddings plays a crucial role in the effectiveness of hate speech detection systems. Our analyses and experiments on the L-HSAB dataset underscore the limitations of relying on pre-trained embeddings that are not tailored to the specific linguistic characteristics of Levantine Arabic.

We evaluated several embedding techniques to assess their performance in detecting hate speech within the L-HSAB dataset. The methods included traditional approaches like Bag-of-Words (BoW) (using unigrams) and Term Frequency-Inverse Document Frequency (TF-IDF), as well as neural embeddings such as pre-trained Arabic fastText, custom-trained Word2Vec on Levantine Arabic data, pre-trained GoogleNews Word2Vec, and pre-trained GloVe embeddings (Harris, 1954; Sparck Jones, 1988; Bojanowski et al., 2016; Mikolov et al., 2013; Pennington et al., 2014).

To provide a more concrete evaluation of the methods discussed, we conducted experiments using two classifiers: Logistic Regression (max\_iter=1000) and a Support Vector Classifier (SVC) with a linear kernel. The results in terms of F1 scores are summarized in Table 1.

Notably, the custom dataset-trained Word2Vec model produced relatively low accuracy scores (0.4429 and 0.3527), which we attribute to the very limited size of our training corpus (approximately

Model	Logistic Regression	SVC
BoW (U)	0.7177	0.7147
TF-IDF	0.6553	0.7217
Custom W2V	0.4429	0.3527
Arabic fastText	0.6823	0.6964
GloVe	0.0606	0.0603
GNews W2V	0.0	0.0

Table 1: F1 scores for various embeddings and classifiers on the L-HSAB dataset. BoW (U) stands for Bag-of-Words (unigrams). Custom W2V refers to a Word2Vec model trained on a small custom Levantine Arabic corpus. Arabic fastText, GloVe, and GNews W2V refer to pre-trained embeddings from Arabic fastText, GloVe, and GoogleNews Word2Vec models respectively.

21,959 words). We anticipate that performance would improve substantially with a larger, more representative Levantine Arabic corpus.

**Effective Techniques:** Our experiments revealed that BoW, TF-IDF, pre-trained Arabic fastText, and custom-trained Word2Vec embeddings significantly outperformed the other methods. These techniques achieved higher F1 scores, indicating better precision and recall in identifying hate speech content. However, it is important to note that the BoW approach, relying solely on unigrams, does not capture contextual relationships between words. As a result, its performance can vary significantly depending on the type and structure of the dataset used. The success of these models can be attributed to their alignment with the linguistic properties of Levantine Arabic, either through their focus on Arabic text or customization to the specific dialect.

**Ineffective Techniques:** In stark contrast, pre-trained embeddings like GoogleNews Word2Vec and GloVe, which are primarily trained on English corpora, scored nearly 0% in F1 metrics. This drastic underperformance highlights a critical issue: models trained predominantly on English data fail to recognize or interpret Arabic text accurately. Consequently, they are ineffective for tasks involving Levantine Arabic hate speech detection.

These findings emphasize the importance of domain-specific adaptations in NLP models. Utilizing embeddings and language models that are trained or fine-tuned on Levantine Arabic data is essential for capturing the unique linguistic features and nuances of the dialect. Relying on generic, pre-trained models not only reduces accuracy but also risks missing or misclassifying hate speech,

thereby undermining the effectiveness of detection systems.

By investing in domain-specific models, researchers and technologists can create more accurate and reliable hate speech detection tools. Such tools will be better equipped to handle the linguistic diversity of Levantine Arabic, ultimately contributing to a safer and more inclusive online environment for speakers of all regional dialects.

#### 4 Ethical Considerations in Hate Speech Detection

The dialectal bias identified above privileges one regional dialect over others, and risk marginalizing communities whose voices are already underrepresented in the digital sphere. There are also ethical concerns beyond issues of data bias. False positives—where non-hate speech is misclassified—can result in the suppression of legitimate cultural expressions, especially in a region where language is tightly bound to identity. A prominent example is the misclassification of the Arabic word "shaheed", meaning "martyr", by social media platforms like Meta ([The Oversight Board, 2024](#)). The term holds significant cultural and religious importance, often used to honor individuals who have died for a sacred cause. However, automated moderation systems have frequently removed content containing "shaheed," interpreting it as a reference to terrorism or violent extremism due to its association with entities on terrorism watchlists.

Conversely, false negatives—where actual hate speech goes undetected—allow harmful narratives to spread unchecked, fueling further violence. For example, derogatory terms or slurs specific to a particular region or group may go unnoticed by moderation systems trained primarily on other dialects or on Modern Standard Arabic. In the context of the Syrian conflict, hate speech containing region-specific pejoratives aimed at certain ethnic or sectarian groups might not be recognized as such by models lacking comprehensive dialectal data. This oversight enables the propagation of inflammatory content that can exacerbate tensions and incite real-world violence.

Technologists and researchers have a responsibility to develop models that not only detect hate speech but do so in a way that respects the linguistic and cultural integrity of Levantine Arabic. Practically, ethical considerations are particularly relevant within a conflict-ridden region like the

Levant where the failure to identify and address hate speech content undermines efforts to promote peace and stability. By incorporating diverse linguistic inputs and cultural insights, developers can create more nuanced models that differentiate between harmful content and legitimate expression, thereby protecting both free speech and community safety.

#### 5 Towards More Culturally Aware Language Technologies

Addressing the challenges of hate speech detection in Levantine Arabic requires practical solutions that consider the language's unique properties. [Bergman and Diab \(2022\)](#) offer valuable guidelines for developing effective and ethically sound NLP tools for underrepresented dialects. By incorporating these recommendations, we can create language technologies that are culturally aware and inclusive, specifically tailored to Levantine Arabic.

##### 5.1 Engaging Local Communities

Engaging local communities is essential for capturing the full spectrum of dialectal variations and cultural contexts within Levantine Arabic. The language's rich diversity necessitates collaboration with native speakers from various regions. Involving annotators and experts who possess both language proficiency and deep understanding of local contexts ensures that the linguistic nuances specific to each dialect are accurately represented ([Radiya-Dixit and Bogen, 2024](#)).

##### 5.2 Rethinking Data Collection and Annotation

To overcome dialectal bias, new data collection and annotation strategies must account for Levantine Arabic's specific properties. Given the significant dialectal variations, stratified sampling techniques are crucial for comprehensively capturing the linguistic landscape ([Bergman and Diab, 2022](#)). Annotation processes should prioritize using annotators proficient in specific regional dialects and familiar with local sociopolitical contexts ([Caliskan et al., 2017](#); [Radiya-Dixit and Bogen, 2024](#)). Researchers must be mindful of potential consequences when collecting data from conflict-affected regions, as certain linguistic features can carry sociopolitical implications. Providing transparent annotation guidelines and support systems for annotators is also critical.

### 5.3 Prioritizing Ethical Design

Developing NLP tools for Levantine Arabic must be grounded in ethical design principles that account for the language’s unique properties. Practitioners should carefully consider the granularity of language divisions within Levantine Arabic and strive for inclusivity without compromising annotation quality (Bergman and Diab, 2022). Providing support systems for annotators is essential, especially given potential exposure to disturbing content in conflict-affected regions. By adopting these strategies, researchers can develop hate speech detection models that are equipped to handle Levantine Arabic’s dialectal diversity and cultural contexts, promoting an inclusive digital environment.

## 6 Conclusion

Detecting hate speech in Levantine Arabic presents unique cultural, linguistic, and ethical challenges due to intricate dialectal variations and biased datasets. This highlights the urgent need for more inclusive NLP approaches. By engaging local communities, reimagining data collection, and embedding ethical considerations into technology design, we can develop tools that effectively identify hate speech while honoring the Levant’s rich linguistic diversity. This paper advocates for renewed cultural sensitivity in NLP applications targeting Levantine Arabic. Addressing sociolinguistic complexities and ethical implications enables us to create tools that serve all speakers, enhance detection accuracy, and promote a more just digital environment throughout the Arab world.

## 7 Limitations

This paper offers a conceptual discussion on the challenges of detecting hate speech in Levantine Arabic. While we provide experimental results to assess the effectiveness of different embedding techniques, the practical impact of our recommendations is still limited by the size and representativeness of our training data. Specifically, the custom dataset-trained Word2Vec model produced relatively lower F1 scores than the pre-trained Arabic fastText model, primarily due to the very limited size of our training corpus (approximately 21,959 words). We anticipate that performance would improve substantially with a larger Levantine Arabic corpus. To this end, we have identified three promising morphologically annotated Levantine corpora—Baladi (Lebanese, ~9.6K tokens), Cur-

ras (Palestinian, ~56K tokens), and Nabra (Syrian, ~60K tokens)—which we plan to combine into a more comprehensive Levantine corpus of approximately 125.6K tokens (Al-Haff et al., 2022; Nayouf et al., 2023). We expect that training our Word2Vec model on this expanded corpus will significantly enhance its performance.

Additionally, while we discuss dialectal variations across Syria, Jordan, Palestine, and Lebanon, the linguistic analysis is not exhaustive, and some regional nuances may not be fully represented. Although we reference frameworks such as the playbook by Bergman and Diab (2022), we do not offer a detailed roadmap for creating inclusive and effective hate speech detection models. Future work should therefore focus on both enriching the training data resources and developing concrete tools to operationalize these recommendations, ensuring more accurate, contextually aware, and inclusive hate speech detection in Levantine Arabic.

## Acknowledgments

We would like to express our sincere gratitude to Dr. Nagham El Karhili at the Global Internet Forum to Counter Terrorism, Professor Stevie Bergman at Brown University, Professor Manar Darwish at Haverford College, and Anna Lacy, Digital Scholarship Librarian at Haverford College, for their insightful feedback and guidance throughout this research. Special thanks to the authors of the L-HSAB dataset—Hala Mulki, Hatem Haddad, Chedi Bechikh Ali, and Halima Alshabani—whose valuable work significantly contributed to this study. We also extend our appreciation to Professor Daniel Ritchie, the Computer Science Department at Brown University, and Google Research for offering and directing the exploreCSR program and for their funding support. We are grateful to the Digital Scholarship team at Haverford College for their assistance and support. Lastly, we thank the Marian E. Koshland Integrated Natural Sciences Center at Haverford College for their funding and support. We are also deeply appreciative of the three anonymous reviewers, whose constructive feedback and suggestions helped us strengthen our work.

## References

Monirah A. Al-Ajlan and Mourad Ykhlef. 2018. [Optimized twitter cyberbullying detection based on deep](#)

- learning. In *2018 21st Saudi Computer Society National Computer Conference (NCC)*, pages 1–5.
- Karim Al-Haff, Mustafa Jarrar, Tymaa Hammouda, and Fadi Zaraket. 2022. *Curras + baladi: Towards a Levantine corpus*. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 769–778, Marseille, France. European Language Resources Association.
- Azalden Alakrot, Liam Murray, and Nikola S. Nikolov. 2018. *Dataset construction for the detection of anti-social behaviour in online communication in arabic*. *Procedia Computer Science*, 142:174–181. Arabic Computational Linguistics.
- Nuha Albadi, Maram Kurdi, and Shivakant Mishra. 2020. Are they our brothers? analysis and detection of religious hate speech in the arabic twittersphere. In *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM '18*, page 69–76. IEEE Press.
- Imran Awan. 2014. *Islamophobia and twitter: A typology of online hate against muslims on social media*. *Policy & Internet*, 6(2):133–150.
- Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2023. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press.
- Solon Barocas and Andrew D. Selbst. 2016. *Big data’s disparate impact*. *California Law Review*, 104(3):671–732.
- Emily Bender. 2019. The benderrule: On naming the languages we study and why it matters. *The Gradient*.
- A. Stevie Bergman and Mona T. Diab. 2022. *Towards responsible natural language annotation for the varieties of arabic*. *Preprint*, arXiv:2203.09597.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606*.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. *Semantics derived automatically from language corpora contain human-like biases*. *Science*, 356(6334):183–186.
- Sergio Andrés Castaño-Pulgarín, Natalia Suárez-Betancur, Luz Magnolia Tilano Vega, and Harvey Mauricio Herrera López. 2021. *Internet, social media and online hate speech. systematic review*. *Aggression and Violent Behavior*, 58:101608.
- Karim El Haff, Mustafa Jarrar, Tymaa Hammouda, and Fadi Zaraket. 2022. *Curras + baladi: Towards a levantine corpus*. *Preprint*, arXiv:2205.09692.
- Zellig S. Harris. 1954. *Distributional structure*.
- Md Saroar Jahan and Mourad Oussalah. 2023. *A systematic review of hate speech automatic detection using natural language processing*. *Neurocomputing*, 546:126232.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. *Efficient estimation of word representations in vector space*. *Preprint*, arXiv:1301.3781.
- Hamdy Mubarak, Kareem Darwish, and Walid Magdy. 2017. *Abusive language detection on Arabic social media*. In *Proceedings of the First Workshop on Abusive Language Online*, pages 52–56, Vancouver, BC, Canada. Association for Computational Linguistics.
- Hala Mulki, Hatem Haddad, Chedi Bechikh Ali, and Halima Alshabani. 2019. *L-HSAB: A Levantine Twitter dataset for hate speech and abusive language*. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 111–118, Florence, Italy. Association for Computational Linguistics.
- Amal Nayouf, Tymaa Hammouda, Mustafa Jarrar, Fadi Zaraket, and Mohamad-Bassam Kurdy. 2023. *Nâbra: Syrian Arabic dialects with morphological annotations*. In *Proceedings of ArabicNLP 2023*, pages 12–23, Singapore (Hybrid). Association for Computational Linguistics.
- Samia Naïm. 2012. *53. Dialects of the Levant*, pages 920–935. De Gruyter Mouton, Berlin, Boston.
- M. Obégi. 1971. *The Phonemic System of a Lebanese Arabic Dialect: (microfilm)*. Theses (Dept. of Modern Languages). National Library of Canada.
- Rasha Omran. 2021. *The letter qaf and conflicting grievances*. <https://www.syria.tv/%D8%AD%D8%B1%D9%81-%D8%A7%D9%84%D9%82%D8%A7%D9%81-%D9%88%D8%A7%D9%84%D9%85%D8%B8%D9%84%D9%88%D9%85%D9%8A%D8%A7%D8%AA-%D8%A7%D9%84%D9%85%D8%AA%D8%B9%D8%A7%D8%B1%D8%B6%D8%A9>.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. *GloVe: Global vectors for word representation*. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Evani Radiya-Dixit and Miranda Bogen. 2024. *Beyond english-centric ai: Lessons on community participation from non-english nlp groups*.
- Ahmad Khalaf Sakarna. 2005. *The linguistic status of the modern jordanian dialects*. *Arabica*, 52(4):522–543.
- Karen Sparck Jones. 1988. *A statistical interpretation of term specificity and its application in retrieval*, page 132–142. Taylor Graham Publishing, GBR.
- The Oversight Board. 2024. *Referring to designated dangerous individuals as “shaheed”*.
- Uğur Ümit Üngör. 2020. *Shabbiha: Paramilitary groups, mass violence and social polarization in homs*. *Violence: An International Journal*, 1(1):59–79.