

An Annotated Corpus of Arabic Tweets for Hate Speech Analysis

Wajdi Zaghoauni

Northwestern University in Qatar
Education City, Doha, Qatar

wajdi.zaghouani@northwestern.edu

Md. Rafiul Biswas

Hamad Bin Khalifa University
Doha, Qatar

mbiswas@hbku.edu.qa

Abstract

Identifying hate speech content in the Arabic language is challenging due to the rich quality of dialectal variations. This study introduces a multilabel hate speech dataset in the Arabic language. We have collected 10,000 Arabic tweets and annotated each tweet, whether it contains offensive content or not. If a text contains offensive content, we further classify it into different hate speech targets such as religion, gender, politics, ethnicity, origin, and others. A text can contain either single or multiple targets. Multiple annotators are involved in the data annotation task. We calculated the inter-annotator agreement, which was reported to be 0.86 for offensive content and 0.71 for multiple hate speech targets. Finally, we evaluated the data annotation task by employing a different transformers-based model in which AraBERTv2 outperformed with a micro-F1 score of 0.7865 and an accuracy of 0.786.

1 Introduction

Identifying hate speech has become a critical area of research due to the rise of online platforms where offensive and harmful content can spread rapidly. The engagement of communities in digital platforms is increasing day by day, and so it increases the risk of individuals being affected by the negativity of hate speech content (Costello and Hawdon, 2020). Therefore, the development of an effective computational tool is required to monitor and mitigate such content. While significant advancements have been made in hate speech detection for English and other widely spoken languages, addressing this issue in the Arabic language remains particularly challenging due to the language's complexity and diversity (Alhazmi et al., 2024a).

Arabic is a rich language with numerous dialectal variations across different regions, making natural language processing (NLP) tasks, including

hate speech detection, more difficult than in languages with more standard forms (Alhazmi et al., 2024b). Moreover, Arabic dialects are highly informal in nature, and they vary significantly from Modern Standard Arabic (MSA), adding layers of complexity when developing NLP models (Abdelsamie et al., 2024; Al-Ibrahim et al., 2023). Many of the available Arabic datasets focus on formal Arabic or a specific dialect, which limits the generalizability of hate speech detection models.

In addition to linguistic diversity, the low-resource nature of Arabic dialects also poses a significant challenge (Abdul-Nabi et al., 2024). Many Arabic NLP tasks suffer from a scarcity of annotated datasets, especially tasks like hate speech detection. While hate speech detection in English has benefited from large-scale datasets and well-developed models, similar efforts in Arabic are comparatively limited (Alkomah and Ma, 2022; Kansok-Dusche et al., 2023). The multilabel nature of hate speech—where a text can target multiple identities such as religion, gender, or ethnicity—adds further complexity, making it essential to develop datasets that capture the breadth of hate speech in different forms (Zaghouani et al., 2024).

This study addresses these challenges by introducing a novel multilabel hate speech dataset in Arabic, specifically designed to encompass multiple hate speech targets such as religion, gender, politics, ethnicity, and origin. By annotating a large corpus of Arabic tweets, we aim to provide a resource that can enable more accurate and nuanced detection of hate speech across different target categories. Our dataset is a contributes to the Arabic NLP field, providing an essential resource for future research on hate speech detection in low-resource and dialectally diverse languages.

1 Novelty: The dataset represents a unique, large-scale, annotated corpus of Arabic tweets specifically designed for multilabel hate speech

analysis. It addresses the paucity of resources in this domain and provides a foundational dataset for the training and evaluation of machine learning models.

2 Availability: The annotated dataset and guidelines are made freely accessible to the research community upon request. The dataset is released under the Creative Commons Attribution 4.0 International (CC BY 4.0) license, which allows for free use, distribution, and adaptation, provided that the original work is properly attributed. This open access ensures that the majority of academic and industry researchers can access and utilize the resource.

3. Utility: The guidelines provide clear instructions on accessing and using the resource. The annotation process, which involves various annotators from different Arab countries, is described in detail, ensuring the representativeness and accuracy of the dataset. The dataset can be employed to train and evaluate various machine learning models of hate speech detection in Arabic social media content (Kakulapati et al., 2023; Xian-Yi et al., 2010; Gaid et al., 2019; Pereira-Kohatsu et al., 2019).

Section 2 describes the related works. Section 3 discusses the data collection strategy and data annotation guidelines. Section 4 discusses the descriptive statistics of the dataset and applied the transformer-based BERT model to evaluate the data annotation. Finally, section 5 discusses the limitation of the works and future scope for improvement.

2 Related Works

Numerous research efforts have been conducted to develop and examine Arabic corpora for various NLP applications. However, some of them focused on the collection of hate speech datasets in Arabic. We will categorize the related works into two categories: i) collection of Arabic hate speech corpus and ii) techniques to detect Arabic hate speech.

Collection of Arabic Hate Speech Corpus: Table 1 describes a recent dataset on Arabic hate speech collections. Mulki et al. 2019 collected Arabic Twitter data between March 2018 and February 2019. They collected 5846 tweets containing three labels: normal, abusive, and hate. The search keywords contain refugee, female, Arabs, Druze, tweets belonging to the post, and replies from politicians, feminists, and activists. (Albadi et al., 2018) collected 9,316 tweets classified into three labels: hateful, abusive, or normal. The search key-

words contain racist, religious, and ideological hate speech, and the search spans from March 2018 to August 2018. (Zaghouni et al., 2024) collected tweets between August 12, 2020, and October 4, 2020. It covers diverse contents focusing on emotion, emotion intensity, offensive content, fact-checking and so on. (Alshaalan and Al-Khalifa, 2020) collected 3075 tweets in 2018 on religion, ethnicity, nationality, and gender. The tweet Id belongs to gulf countries. Khezzar et al. 2023 proposed the arHateDataset dataset, which comprises 34,107 tweets obtained from different Arabic standards and dialects. However, this dataset is not a unique corpus from the author. Rather, they combined seven Arabic datasets together to form this dataset.

Techniques to Detect Arabic Hate Speech: Table 1 demonstrates the uses of different model to detect offensive content. (Mulki et al., 2019) used traditional machine learning techniques Support Vector Machine (SVM) and Naïve Bayes (NB) to identify tweets whether a tweet is hateful or not. (Albadi et al., 2018) applied SVM and Logistic Regression (LR) techniques as traditional machine learning model. For deep learning model they used CNN, RNN, GRU, CNN+GRU and for transformer, they applied BERT model. CNN outperformed on all of these experiments. (Zaghouni et al., 2024) applied different machine learning techniques such as LR, SVM, Random Forest (RF) and transformer based model AraBERT to classify the content either offensive or not. AraBERT scored for micro F1 score 0.65. (Alsafari et al., 2020) fused transformer BERT base model with machine learning model (LR and SVM) and with deep learning model (CNN, GRU, LSTM). The fusion model mBERT+CNN outperformed (micro F1 Score 0.7899) over other techniques. (Khezzar et al., 2023) applied various techniques for hate speech classification and AraBERT obtained best accuracy score (0.91) among them.

These studies highlight the importance of developing annotated Arabic corpora for reliable hate speech detection within the challenges of informal dialectal social media posts. To extend the existing Arabic corpus, our dataset provides multilabel hate speech which is unique. Annotating each tweet to all the possible labels has been proposed for the first time in our dataset. Moreover, we experiments the performance of transformer based model in identifying the annotated label.

Dataset	Keywords	Multi-label	Size	Platforms	Dialect	Model
(Mulki et al., 2019)	Refugee, female, Arabs	No	5,846	Twitter	Syrian, Lebanese	SVM, NB
(Albadri et al., 2018)	Racist, religious, tribes	No	9,316	Twitter	Saudi	SVM, CNN, RNN, BERT
(Zaghouani et al., 2024)	Variety of topics	Yes	15,965	Twitter	Across Arab countries	LR, SVM, AraBERT
(Alsaifari et al., 2020)	Religion, ethnicity, gender	No	3,075	Twitter	Gulf countries	LR, CNN, mBERT
(Khezzar et al., 2023)	Varieties	No	34,107	Twitter	Multi	NB, CNN, AraBERT

Table 1: Comparison of Arabic Hate Speech Datasets

3 Methodology

3.1 Data Collection

We collected 60 million Arabic tweets from August 12 and October 4, 2020, and randomly sampled 600,000 tweets (1% of the original). After removing duplicates, incomplete tweets, short tweets (less than 5 words), and long tweets (more than 80 words), we obtained a diverse set of tweets for annotation. In this current study, we used ASAD tools (Hassan et al., 2021) to classify offensive content from raw tweets. ASAD is a suite of tools for analyzing Arabic social media content, primarily focused on Twitter. It offers modules for tasks such as Dialect Identification, Sentiment and Emotion Analysis, Offensive Language and Hate Speech Detection, Spam Detection, and more. We identified offensive content by choosing 4000 tweets from the highest confidence (80%-100%), 4000 tweets from average confidence (60%-79%), and 2000 tweets from low confidence (1%-39%). Then, we applied sentiment classification and randomly selected 4000 tweets from positive, 4000 tweets from negative, and 2000 from neutral categories. Next, we selected tweets that were marked as hate speech by ASAD tools. We also included tweets that were labeled as adult content and spam. Figure 1 shows the data collection process.

It is noteworthy that there exists a degree of overlap among the extracted tweets from the aforementioned selection steps. For instance, a tweet might simultaneously contain offensive content and adult content. We removed duplicates while merging the tweets and were left with 10,000 tweets. For this study, we focused on emotions and multilevel classification of offensive/hate speech content.

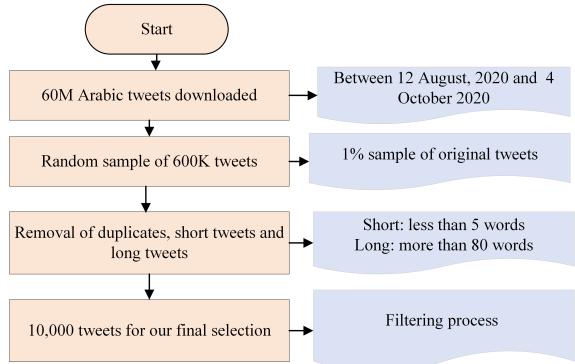


Figure 1: Filtering dataset for research

3.2 Ethical Approval

The authors carefully considered the ethical implications of the resource and applied for and received an institutional IRB exemption. The dataset was collected in compliance with Twitter (X) API rules, and user privacy was protected through the anonymization of personally identifiable information. The study adheres to established ethical guidelines for social media research.

3.3 Annotation Procedure

The annotation task was performed by several annotators (ranging from one to five based on the understanding of the dialects) and a manager (who led the project).

Annotators Recruiting and Training: The annotators were recruited from different Arab countries, including two from Tunisia, one from Egypt, one from Jordan, and one from Qatar. Among them, three were female, and two were male. This diverse representation ensures that the dataset accurately captures the linguistic diversity and cultural contexts within the Arab-speaking world. Annotators from different countries provide insight into how certain words or phrases might be interpreted dif-

ferently across cultures, improving the quality and accuracy of annotations. The annotators underwent a training process to ensure high-quality and consistent annotations. Each category had practice questions to familiarize them with the types of content they would encounter. Annotators met regularly to discuss difficulties and revise their approaches as needed.

Role of the Annotation Manager: The annotation manager evaluated the annotation process. He resolved if there were any conflicts among annotators and chose the final label. The manager also analyzed errors, calculated inter-annotator agreement, provided feedback, and updated the guidelines to maximize the consistent label.

Annotation Platform: Annotation was performed using MicroMappers, an online tool to manage annotation. Annotators worked individually using this tool but met weekly to discuss their outcomes and difficulties.

3.4 Annotation Guidelines

Proper annotation guidelines are essential to building an appropriate dataset with fewer errors (Zaghouni et al., 2014). Also, a proper guideline can minimize subjective bias among annotators (Röttger et al., 2021). Figure 2 depicts a glimpse of annotation guidelines. This annotation is a multilabel procedure.

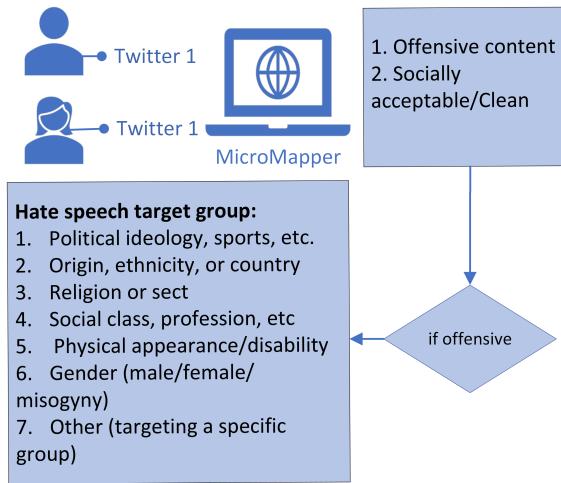


Figure 2: Data annotation procedure

Offensive: Offensive content refers to cause harm, insult, or discomfort to individuals or groups. Example: "This is a lie from you and has no source except from a seller among you".

Socially acceptable\Clean: Tweet content does not carry offensive meaning and is socially accept-

able to all.

Example: "And old customers, there is nothing left for them.. A failed marketing method. Your customer who has been with you for years, discount the value of the invoice."

Hate Speech Target: We distinguished offensive content as hate speech when targeted at a specific group. We further labeled hateful content into seven categories. Hate speech targeting: (1) Political ideology, sports, etc. (e.g., "Liberals are destroying our country!"); (2) Origin, ethnicity, or country (e.g., "Go back to your country!"); (3) Religion or sect (e.g., "All Muslims are terrorists"); (4) Social class, profession, etc. (e.g., "Homeless people are lazy"); (5) Physical appearance/disability/rich/poor (e.g., "Fat people are disgusting"); (6) Gender (male/female/misogyny) (e.g., "Women belong in the kitchen") and (7) Other (targeting a specific group) (e.g., "LGBTQ+ people are unnatural")."

Some text target more than one group. For example, "She looks like a Bangladeshi who works in the grocery store in the neighborhood." This tweet bears offensive content. It has targeted two groups: "origin, ethnicity, or country" and "physical appearance/disability". While annotating each tweet, we consider all the possible target group. Thus, if multiple annotators annotate one tweet offensive and then they have diverse opinion in target group, we combine all the target group. This introduces the data annotation diversity according to dialectal variation.

4 Results and Discussion

4.1 Offensiveness Analysis

Table 2 presents the frequency of offensiveness data. The majority of tweets (60.36%) were identified as hateful or offensive, highlighting a significant presence of negative content in the dataset. Clean or socially acceptable tweets counted for (37.19%), showing that a substantial portion of the content was non-offensive. There are some tweets that do not fall into the two categories. They belong to different meanings and are very few in number. For instance, a small portion of the tweets are labeled as vulgar (0.018%), while others contain violent content (0.006%). Additionally, a few tweets exhibit irony or sarcasm (0.002%) and humor (0.002%), highlighting the rarity of these categories.

Offensiveness	N	%
Hateful/Offensive	6036	60.36
Clean/ Socially Acceptable	3719	37.19
Violence	63	.006
Irony/sarcasm	26	.002
Humor	20	.002
Vulgar	184	.018

Table 2: Frequency of offensiveness data (N represents number of items)

4.2 Hate Speech Target

As discussed before, we further classified the offensive content (6036 tweets) into hate speech target group. Table 3 provides descriptive data for hate speech targets which is categorized into seven distinct labels. The most common target was political ideology, sports, etc., accounting for (28.76%). The "Other" category, encompassing a specific group, was almost as prevalent, with (28.51%) of the tweets. Hate speech targeting origin, ethnicity, or country was also notable, representing (26.20%) of the tweets, highlighting significant issues related to racism and xenophobia. Religion or sect was the target in (13.51%) of the tweets, reflecting religious intolerance. Hate speech aimed at social class, profession, etc., appeared in (2.47%) of the tweets. Physical appearance, disability, wealth status, and gender were less frequent targets, with (0.25%) and (0.30%), respectively. Table 3 shows the hate speech target results.

Hate Speech Target	N	%
Political ideology, sports	1,992	28.8
Other	1,975	28.5
Origin, ethnicity, country	1,815	26.2
Religion or sect	936	13.5
Social class, profession	171	2.5
Physical appearance, disability	17	0.3
Gender	21	0.3

Table 3: Frequency of hate speech target

4.3 Inter Annotator Agreement

To measure inter annotator agreement, we took a sample of 500 tweets and asked all the annotators to annotate the tweets into two categories: offensive or not (we didn't measure for hate speech target as

we combine all possible annotation label to maximize the target group). As there were more than two annotators involved in this procedure, we applied Fleiss' Kappa to calculate Inter Annotator Agreement (IAA).

$$\kappa = \frac{P_o - P_e}{1 - P_e} \quad (1)$$

where:

$$P_o = \frac{1}{N \cdot (N-1)} \sum_{i=1}^N \left[\sum_{j=1}^N \left(P_{ij}^2 - \frac{P_i \cdot P_j}{(N-1)} \right) \right] \quad (2)$$

and

$$P_e = \frac{1}{N \cdot (N-1)} \sum_{i=1}^N \left(\sum_{j=1}^N \left(\frac{P_i \cdot P_j}{N} \right) \right) \quad (3)$$

where N is the number of raters, P_{ij} is the proportion of times that the i -th rater assigned a particular category j , and P_i and P_j are the proportions of times categories i and j were assigned, respectively.

We achieved an Inter-Annotator Agreement (IAA) score of 0.8143 among the annotators, which indicates a high level of agreement among them. A Fleiss' Kappa value of 0.8143 suggests that the agreement among the annotators is substantial, reflecting a high consistency in their evaluations.

4.4 Annotation Evaluation with Transformer Model

To evaluate the data annotation task, we implemented transformer based model to classify tweet into offensive or not. We also applied transformer model for hate speech target classification. We chose transformer model because it has shown better performance in hate speech classification (Alsaifari et al., 2020; Albadi et al., 2018; Khezzar et al., 2023).

Model Architecture: Firstly, we prepossessed the text by removing link, username, empty spaces, emojis and weird characters. Then, we applied one hot encoding to each label. The dataset was split into 70% training, 20% testing, and 10% validation. We used the training dataset to finetune the pre-trained BERT model. We validate the model outcomes through validation dataset. Finally, we used test dataset to predict the classification layer.

The model was set up with batch=4, epoch=3, learning rate=2e-5, loss function:sigmoid. We tried with various BERT-base models. Initially, we fine-tuned the newest version of AraBERT model which is ‘aubmindlab/bert-base-arabertv2’ (Antoun et al.). Then we fine-tuned with pre-trained CAMeLBERT (Inoue et al., 2021) model. We also fine-tuned with XLM-RoBERTa (Conneau, 2019).

Model Evaluation: Table 4 shows the performance metrics for different classifications. Offensive Vs Clean is binary classification problem. The models are quite good to accurately predict the outcomes. In this cases, AraBERT pre-trained model outperformed with F1-score of 0.7865 and accuracy of 0.7860. CamelBERT also achieved good score 0.7360.

For hate speech target classification, AraBERTv2 outperforms in all metrics with F1-score 0.6889, Accuracy 0.5985, Precision 0.7172 and Recall 0.6628. It provides a better balance between precision and recall, making it more effective at identifying hate speech targets overall.

Offensive Vs Clean				
Model	F1-Score	Accuracy	Precision	Recall
AraBERT	0.7865	0.7860	0.7865	0.7865
CamelBERT	0.7360	0.7355	0.7360	0.7360
xlm-roberta	0.6280	0.6280	0.6280	0.6280
Hate Speech Target				
xlm-roberta	0.6607	0.5771	0.7098	0.6181
AraBERTv2	0.6889	0.5985	0.7172	0.6628
CamelBERT	0.6814	0.6002	0.7234	0.6440

Table 4: Performance metrics for Offensive vs Clean and Hate Speech Target classification

System Architecture: The training was conducted on a Google Colab Server with NVIDIA Tesla T4 GPU with 16GB of VRAM. The experiment was run several times with hyperparameter tuning like epoch, batch size, learning rate, weighted decay and so on. We saved the best weight by comparing the Micro F1-score and predicted the output with the best weight model.

5 Limitation and Future Scope

Our hate speech dataset collection is not beyond limitation.

- Dataset Size:** Our dataset contains only 10000 tweets which is limited for generability

of hate speech detection. While every day people posts a lots in the social media, we need to increase the volume of the dataset.

- Costly:** X data (Former Twitter data) was freely accessible to the research communities. Now everyone has to pay to use the X API which is costly to download data and put limitation on doing research on X data.
- Dialectal variation:** Even though our dataset contains data from all over region, we didn’t distuinsh the data based on dialectal region. It reduces the model performance while predicting the outcomes.

6 Conclusion

Our dataset addresses a critical gaps in low resources language, particularly due to the linguistic diversity and dialectal variations in Arabic. Through careful annotation procedures and high inter-annotator agreement, we demonstrate the dataset’s reliability and effectiveness. Evaluation using advanced transformer-based models like AraBERT further highlights its practical utility and robust performance.

Data and Code Availability

The dataset is released under the CC BY 4.0 license, allowing free use, distribution, and adaptation with proper attribution. The user can freely access this dataset by filling in a consent form: <https://forms.gle/S9fZtYjAyLAqFsH19>. Code can be accessed by github.com/rafiulbiswas/HateSpeechClassification.

Acknowledgments

This study was partially funded by the grant NRP14C-0916-210015, awarded by the Qatar Research, Development and Innovation Council (QRDI).

References

Mahmoud Mohamed Abdelsamie, Shahira Shaaban Azab, and Hesham A Hefny. 2024. A comprehensive review on arabic offensive language and hate speech detection on social media: methods, challenges and solutions. *Social Network Analysis and Mining*, 14(1):1–49.

Razan Abdul-Nabi, Rasha Obeidat, and Anas Bsoul. 2024. A survey on machine translation of low-resource arabic dialects. In *2024 15th International*

Conference on Information and Communication Systems (ICICS), pages 1–6. IEEE.

Rogayah M Al-Ibrahim, Mostafa Z Ali, and Hassan M Najadat. 2023. Detection of hateful social media content for arabic language. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 22(9):1–26.

Nuha Albadi, Maram Kurdi, and Shivakant Mishra. 2018. 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*, pages 69–76. ACM.

Ali Alhazmi, Rohana Mahmud, Norisma Idris, Mohamed Elhag Mohamed Abo, and Christopher Eke. 2024a. A systematic literature review of hate speech identification on arabic twitter data: research challenges and future directions. *PeerJ Computer Science*, 10:e1966.

Ali Alhazmi, Rohana Mahmud, Norisma Idris, Mohamed Elhag Mohamed Abo, and Christopher Ifeanyi Eke. 2024b. Code-mixing unveiled: Enhancing the hate speech detection in arabic dialect tweets using machine learning models. *Plos one*, 19(7):e0305657.

Fatimah Alkomah and Xiaogang Ma. 2022. A literature review of textual hate speech detection methods and datasets. *Information*, 13(6):273.

Safa Alsaafari, Samira Sadaoui, and Malek Mouhoub. 2020. Hate and offensive speech detection on arabic social media. *Online Social Networks and Media*, 19:100096.

Raghad Alshaalan and Hend Al-Khalifa. 2020. Hate speech detection in saudi twittersphere: A deep learning approach. In *Proceedings of the fifth Arabic natural language processing workshop*, pages 12–23.

Wissam Antoun, Fady Baly, and Hazem Hajj. Arabert: Transformer-based model for arabic language understanding. In *LREC 2020 Workshop Language Resources and Evaluation Conference 11–16 May 2020*, page 9.

A Conneau. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

Matthew Costello and James Hawdon. 2020. Hate speech in online spaces. *The Palgrave handbook of international cybercrime and cyberdeviance*, pages 1397–1416.

Bharat Gaind, Varun Syal, and Sneha Padgalwar. 2019. Emotion detection and analysis on social media. *arXiv preprint arXiv:1901.08458*.

Sabit Hassan, Hamdy Mubarak, Ahmed Abdelali, and Kareem Darwish. 2021. Asad: Arabic social media analytics and understanding. In *Proceedings of the 16th Conference of the European Chapter of the*

Association for Computational Linguistics: System Demonstrations, pages 113–118.

Go Inoue, Bashar Alhafni, Nurpeis Baimukan, Houda Bouamor, and Nizar Habash. 2021. The interplay of variant, size, and task type in Arabic pre-trained language models. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, Kyiv, Ukraine (Online). Association for Computational Linguistics.

V Kakulapati, Sheri Mahender Reddy, and A Paramasivam. 2023. Managing postpandemic effects using artificial intelligence with human-computer interaction. In *Innovations in Artificial Intelligence and Human-Computer Interaction in the Digital Era*, pages 207–232. Elsevier.

Julia Kansok-Dusche, Cindy Ballaschk, Norman Krause, Anke Zeißig, Lisanne Seemann-Herz, Sebastian Wachs, and Ludwig Bilz. 2023. A systematic review on hate speech among children and adolescents: Definitions, prevalence, and overlap with related phenomena. *Trauma, violence, & abuse*, 24(4):2598–2615.

Ramzi Khezzar, Abdelrahman Moursi, and Zaher Al Aghbari. 2023. arhatedetector: detection of hate speech from standard and dialectal arabic tweets. *Discover Internet of Things*, 3(1):1.

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.

Juan Carlos Pereira-Kohatsu, Lara Quijano-Sánchez, Federico Liberatore, and Miguel Camacho-Collados. 2019. Detecting and monitoring hate speech in twitter. *Sensors*, 19(21):4654.

Paul Röttger, Bertie Vidgen, Dirk Hovy, and Janet B Pierrehumbert. 2021. Two contrasting data annotation paradigms for subjective nlp tasks. *arXiv preprint arXiv:2112.07475*.

Cheng Xian-Yi, Zhu Ling-ling, Zhu Qian, and Wang Jin. 2010. The framework of network public opinion monitoring and analyzing system based on semantic content identification. *Journal of Convergence Information Technology*, 5(10):1–5.

Wajdi Zaghouani, Behrang Mohit, Nizar Habash, Osama Obeid, Nadi Tomeh, Alla Rozovskaya, Noura Farra, Sarah Alkuhlani, and Kemal Oflazer. 2014. Large scale arabic error annotation: Guidelines and framework.

Wajdi Zaghouani, Hamdy Mubarak, and Md Rafiul Biswas. 2024. So hateful! building a multi-label hate speech annotated arabic dataset. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15044–15055.