

# Arabic Offensive Language on Twitter: Analysis and Experiments

Hamdy Mubarak<sup>1</sup> Ammar Rashed<sup>2</sup> Kareem Darwish<sup>1</sup>  
Younes Samih<sup>1</sup> Ahmed Abdelali<sup>1</sup>

<sup>1</sup>Qatar Computing Research Institute, HBKU

<sup>2</sup>Özyeğin University

{hmubarak, kdarwish, ysamih, aabdelali}@hbku.edu.qa  
ammar.rasid@ozu.edu.tr

## Abstract

Detecting offensive language on Twitter has many applications ranging from detecting/predicting bullying to measuring polarization. In this paper, we focus on building a large Arabic offensive tweet dataset. We introduce a method for building a dataset that is not biased by topic, dialect, or target. We produce the largest Arabic dataset to date with special tags for vulgarity and hate speech. We thoroughly analyze the dataset to determine which topics, dialects, and gender are most associated with offensive tweets and how Arabic speakers use offensive language. Lastly, we conduct many experiments to produce strong results ( $F1 = 83.2$ ) on the dataset using SOTA techniques.

## 1 Introduction

**Disclaimer:** Due to the nature of the paper, some examples herein contain highly offensive language and hate speech. They don't reflect the views of the authors in any way. This work is an attempt to help fight such speech.

Much recent interest has focused on the detection of offensive language and hate speech in online social media. Offensiveness is often associated with undesirable behaviors such as trolling, cyberbullying, online extremism, political polarization, and propaganda. Thus, offensive language detection is instrumental for a variety of application such as: quantifying polarization (Barberá and Sood, 2015; Conover et al., 2011), trolls and propaganda account detection (Darwish et al., 2017), hate crimes likelihood estimation (Waseem and Hovy, 2016); and predicting conflicts (Chadefaux, 2014). In this paper, we describe our methodology for building a large dataset of Arabic offensive tweets. Given that roughly 1-2% of all Arabic tweets are offensive (Mubarak and Darwish, 2019), targeted annotation is essential to efficiently build

a large dataset. Since our methodology does not use a seed list of offensive words, it is not biased by topic, target, or dialect. Using our methodology, we tagged a 10,000 Arabic tweet dataset for offensiveness, where offensive tweets account for roughly 19% of the tweets. Further, we labeled tweets as vulgar or hate speech. To date, this is the largest available dataset, which we plan to make publicly available along with annotation guidelines. We use this dataset to characterize Arabic offensive language to ascertain the topics, dialects, and users' gender that are most associated with the use of offensive language. Though we suspect that there are common features that span different languages and cultures, some characteristics of Arabic offensive language are language and culture specific. Thus, we conduct a thorough analysis of how Arab users use offensive language. Next, we use the dataset to train strong Arabic offensive language classifiers using state-of-the-art representations and classification techniques. Specifically, we experiment with static and contextualized embeddings for representation along with a variety of classifiers such as Transformer-based and Support Vector Machine (SVM) classifiers. The contributions of this paper are as follows:

- We built the largest Arabic offensive language dataset to date that is also labeled for vulgar language and hate speech and is not biased by topic or dialect. We describe the methodology for building it along with annotation guidelines.
- We performed thorough analysis to describe the peculiarities of Arabic offensive language.
- We experimented with SOTA classification techniques to provide strong results on detecting offensive language.

## 2 Related Work

Many recent papers have focused on the detection of offensive language, including hate speech (Agrawal and Awekar, 2018; Badjatiya et al., 2017; Davidson et al., 2017; Djuric et al., 2015; Kwok and Wang, 2013; Malmasi and Zampieri, 2017; Nobata et al., 2016; Yin et al., 2009). Offensive language can be categorized as: *vulgar*, which include explicit and rude sexual references, *pornographic*, and *hateful*, which includes offensive remarks concerning people’s race, religion, country, etc. (Jay and Janschewitz, 2008). Prior works have concentrated on building annotated corpora and training classification models. Concerning corpora, [hatespeechdata.com](http://hatespeechdata.com) attempts to maintain an updated list of hate speech corpora for multiple languages including Arabic and English. Further, SemEval 2019 ran an evaluation task targeted at detecting offensive language, which focused exclusively on English (Zampieri et al., 2019). For SemEval 2020, they extended the task to include other languages including Arabic (Zampieri et al., 2020). As for classification models, most studies used supervised classification at either word level (Kwok and Wang, 2013), character sequence level (Malmasi and Zampieri, 2017), and word embeddings (Djuric et al., 2015). The studies used different classification techniques including Naïve Bayes (Kwok and Wang, 2013), Support Vector Machines (SVM) (Malmasi and Zampieri, 2017), and deep learning (Agrawal and Awekar, 2018; Badjatiya et al., 2017; Nobata et al., 2016) classification. The accuracy of the aforementioned system ranged between 76% and 90%. Earlier work looked at the use of sentiment words as features as well as contextual features (Yin et al., 2009).

The work on Arabic offensive language detection is relatively nascent (Abozinadah, 2017; Alakrot et al., 2018; Albadi et al., 2018; Mubarak et al., 2017; Mubarak and Darwish, 2019). Mubarak et al. (2017) suggested that certain users are more likely to use offensive languages than others, and they used this insight to build a list of offensive Arabic words and to construct a labeled set of 1,100 tweets. Abozinadah (2017) used supervised classification based on a variety of features including user profile features, textual features, and network features. They reported an accuracy of nearly 90%. Alakrot et al. (2018) used supervised classification based on word n-grams to detect offensive language in YouTube comments. They im-

proved classification with stemming and achieved a precision of 88%. Albadi et al. (2018) focused on detecting religious hate speech using a recurrent neural network.

Arabic is a morphologically rich language with a standard variety called Modern Standard Arabic (MSA), which is typically used in formal communication, and many dialectal varieties that differ from MSA in lexical selection, morphology, phonology, and syntactic structures. In MSA, words are typically derived from a set of thousands of roots by fitting a root into a stem template and the resulting stem may accept a variety of prefixes and suffixes. Though word segmentation, which greatly improves word matching, is quite accurate for MSA (Abdelali et al., 2016), with accuracy approaching 99%, dialectal segmentation is not sufficiently reliable, with accuracy ranging between 91-95% for different dialects (Samih et al., 2017). Since dialectal Arabic is ubiquitous in Arabic tweets and many tweets have creative spellings of words, recent work on Arabic offensive language detection used character-level models (Mubarak and Darwish, 2019).

## 3 Data Collection

### 3.1 Collecting Arabic Offensive Tweets

Our target is to build a large Arabic offensive language dataset that is representative of its appearance on Twitter and is hopefully not biased to specific dialects, topics, or targets. One of the main challenges is that offensive tweets constitute a very small portion of overall tweets. To quantify their proportion, we took 3 random samples of tweets from different days, with each sample composed of 1,000 tweets, and we found that only 1-2% of them were offensive (including pornographic advertisements). This percentage is consistent with previously reported percentages (Mubarak et al., 2017). Thus, annotating random tweets is grossly inefficient. One way to overcome this problem is to use a seed list of offensive words to filter tweets. However, doing so is problematic, as it would skew the dataset to particular types of offensive language or to specific dialects. Offensiveness is often dialect and country specific.

After inspecting many tweets, we observed that many offensive tweets have the vocative particle  $\text{يَا}$  (“yA” – meaning “O”)<sup>1</sup>, which is mainly used

---

<sup>1</sup>Arabic words are provided along with their Buckwalter

in directing the speech to a specific person or group. The ratio of offensive tweets increases to 5% if a tweet contains one vocative particle and to 19% if it has at least two vocative particles. Users often repeat this particle for emphasis, as in: يا أمي يا حنونة (“yA Amy yA Hnwnp” – O my mother, O kind one), which is endearing and non-offensive, and يا كلب يا قذر (“yA klb yA q\*r” – “O dog, O dirty one”), which is offensive. We decided to use this pattern to increase our chances of finding offensive tweets. One of the main advantages of the pattern يا ... يا (“yA ... yA”) is that it is not associated with any specific topic or genre, and it appears in all Arabic dialects. Though the use of offensive language does not necessitate the appearance of the vocative particle, the particle does not favor any specific offensive expressions and greatly improves our chances of finding offensive tweets. Using Twitter APIs, we collected 660k Arabic tweets having this pattern between April 15 – May 6, 2019. To increase diversity, we sorted the word sequences between the vocative particles and took the most frequent 10,000 unique sequences. For each word sequence, we took a random tweet containing that sequence. Then we annotated those tweets, ending up with 1,915 offensive tweets which represent roughly 19% of all tweets. Each tweet was labeled as: offensive, which could additionally be labeled as vulgar and/or hate speech, or Clean. We describe in greater detail our annotation guidelines, which are compatible with the OffensEval2019 annotation guidelines (Zampieri et al., 2019). For example, if a tweet has insults or threats targeting a group based on their nationality, ethnicity, gender, political affiliation, religious belief, or other common characteristics, this is considered hate speech (Zampieri et al., 2019). It is worth mentioning that we also considered insulting groups based on their sport affiliation as a form of hate speech. Often, being a fan of a particular sporting club is considered a part of the personality that rarely changes over time (similar to religious and political affiliations). Many incidents of violence have occurred among fans of rival clubs.

Although we used a generic pattern that is used across dialects and topics, such may not cover all the stylistic diversity of offensive expressions. However, our approach considerably narrows the search space for offensive tweets, which consti-

transliteration and English translation.

tute a small percentage of tweets in general, while being far more generic than using a seed list of offensive words, which may greatly skew the distribution of offensive tweets. For future work, we plan to explore other methods for identifying offensive tweets with greater stylistic diversity.

### 3.2 Annotating Tweets

We developed annotation guidelines jointly with an experienced annotator, who is a native Arabic speaker with good knowledge of various Arabic dialects, in accordance to the OffensEval2019 guidelines. Tweets were given one or more of the following four labels: *offensive*, *vulgar*, *hate speech*, or *clean*. Since the *offensive* label covers both vulgar and hate speech and vulgarity and hate speech are not mutually exclusive, a tweet can be just offensive or offensive and vulgar and/or hate speech. The annotation adhered to the following guidelines:

**OFFENSIVE (OFF):** Offensive tweets contain explicit or implicit insults or attacks against other people, or inappropriate language, such as:

**Direct threats or incitement**, ex: احرقوا مقرات المعارضة (“Ahrqwa mqrAt AlmEArDp” – “burn opposition headquarters”) and اقتلوا هذا المنافق (“h\*A AlmnAfq yjb qtlh” – “kill this hypocrite”).

**Insults and expressions of contempt**, which include: *Animal analogies*, ex: يا كلب (“yA klb” – “O dog”) and كل تبين (“kl tbn” – “eat hay”);

*Insult to family*, ex: يا روح أمك (“yA rwH Amk” – “O mother’s soul”); *Sexually-related insults*, ex:

يا ديوث (“yA dywv” – “O cuckold”); *Damnation*,

ex: الله يلعنك (“Allh ylenk” – “may God curse you”); and *Attacks on morals*, ex:

يا كاذب (“yA kA\*b” – “O liar”).

**VULGAR (VLG):** Vulgar tweets are offensive tweets that contain profanity, such as mentions of private parts or sexual-related acts or references.

**HATE SPEECH (HS):** Hate speech tweets are offensive tweets targeting group based on common characteristics such as: *Race*, ex: يا زنجي (“yA znjy” – “O Negro”);

*Ethnicity*, ex. الفرس الأنجاس (“Alfrs AlAnjAs” – “Impure Persians”); *Group or party*, ex: أبوك شيوعي (“Abwk \$ywy” – “your father is a communist”); and *Religion*, ex: دينك القذر

دينك القذر (“Dink alqdr” – “your religion is dirty”).

(“dynk Alq\*r” – “your filthy religion”).

**CLEAN (CLN):** Clean tweets do not contain vulgar or offensive language. We noticed that some tweets have some offensive words, but the whole tweet should not be considered as offensive due to the intention of users. This suggests that normal string match without considering contexts may fail in some cases. Examples of such ambiguous cases include: *Humor*, ex: `يا عدوة الفرحة ههه` (“yA Edwp AlfrHp hhh” – “O enemy of happiness hahaha”); *Advice*, ex: `لا تقل لصاحبك يا خنزير` (“1A tql lSAHbk yA xnzyr” – “don’t say to your friend: You are a pig”); *Condition*, ex: `إذا عارضتهم يقولون يا عميل` (“A\*A EArdthm yqwlwn yA Emyl” – “if you disagree with them, they call you a spy”); *Condemnation*, ex: `لماذا نسب بقول: يا بقرة؟` (“lMA\*A nsb bqwl: yA bqrp?” – “Why do we insult others by saying: O cow?”); *Self offense*, ex: `تعبت من لساني القذر` (“tEbt mn lsAny Alq\*r” – “I am tired of my dirty tongue”); *Non-human target*, ex: `يا بنت المجنونة يا كورة` (“yA bnt Almjnwnp yA kwrp” – “O daughter of the crazy one O football”); and *Quotation from a movies or a story*, ex: `تاني يا زكي! تاني يا فاشل` (“tAny yA zky! tAny yA fA\$l” – “again smarty! again O loser”). For ambiguous expressions, the annotator searched Twitter to observe real sample usages.

Table 1 shows the distribution of the annotated tweets. There are 1,915 offensive tweets, including 225 vulgar tweets and 506 hate speech tweets, and 8,085 clean tweets. To validate annotation quality, we asked three additional annotators to annotate two tweet sample sets. The first was a random sample of 100 tweets containing 50 offensive and 50 non-offensive tweets. The Inter-Annotator Agreement (IIA) between the annotators using Fleiss’s Kappa coefficient (Fleiss, 1971) was 0.92. The second was general random samples containing 100 tweets each from the dataset, and the IIA with the dataset was: 0.97, 0.96, and 0.97. This high level of agreement gives more confidence in the quality of the annotation. Data can be downloaded from:

<https://alt.qcri.org/resources/OSACT2020-sharedTask-CodaLab-Train-Dev-Test.zip>

	Tweets	Words
Offensive	1,915	38k
– Vulgar	225	4k
– Hate speech	506	13k
Clean	8,085	151k
Total	10,000	193k

Table 1: Distribution of offensive and clean tweets.

### 3.3 Statistics and User Demographics

Given the annotated tweets, we wanted to ascertain the distribution of: types of offensive language, genres or topics where it is used, the dialects used, and the gender of users using such language. Accordingly, the annotator manually examined and tagged all the offensive tweets.

**Topic:** Figure 1 shows the distribution of topics associated with offensive tweets. As the figure shows, sports and politics are most dominant for offensive language including vulgar and hate speech.

**Dialect:** We looked at MSA and four major dialects, namely Egyptian (EGY), Leventine (LEV), Maghrebi (MGR), and Gulf (GLF). Figure 2 shows that 71% of vulgar tweets were written in EGY followed by GLF, which accounted for 13% of vulgar tweets. MSA was not used in any vulgar tweets. As for offensive tweets in general, EGY and GLF were used in 36% and 35% of the offensive tweets respectively. Unlike the case of vulgar language, 15% of the offensive tweets were written in MSA. For hate speech, GLF and EGY were again dominant and MSA constituted 21% of the tweets. This is consistent with findings for other languages, e.g. English and Italian, where vulgarity was more frequently associated with colloquial language (Mattiello, 2005; Maisto et al., 2017).

**Gender:** Figure 3 shows that the vast majority of offensive tweets, including vulgar and hate speech, were authored by males. Female Twitter users accounted for 14% of offensive tweets in general and 6% and 9% of vulgar and hate speech respectively. Figure 4 shows a detailed categorization of hate speech types, where the top three include insulting groups based on their political ideology, origin, and sport affiliation. Religious hate speech appeared in only 15% of all hate speech tweets.

Next, we analyzed all tweets labeled as offensive to better understand how Arabic speakers use offensive language. Here is a breakdown of usage:

**Direct name calling:** The most frequent attack is to call a person an animal name, and the most used




Figure 1: Topic distribution for offensive language and its sub-categories




Figure 2: Dialect distribution for offensive language and its sub-categories

*animals* were كلب (“klb” – “dog”), حمار (“HmAr” – “donkey”), and بهيم (“bhym” – “beast”). The second most common was insulting mental abilities using words such as غبي (“gby” – “stupid”) and عيبط (“EbyT” – “idiot”). Culturally, not all animal names are used as insults. For example, animals such as أسد (“Asd” – “lion”), صقر (“Sqr” – “falcon”), and غزال (“gzAl” – “gazelle”) are typically used for praise. For other insults, people use: some *bird* names such as دجاجة (“djAjp” – “chicken”), بومة (“bwmp” – “owl”), and غراب (“grAb” – “crow”); *insects* such as ذبابة (“\*bAbp” – “fly”), صرصور (“SrSwr” – “cockroach”), and حشرة (“H\$rp” – “insect”); *microorganisms* such as جرثومة (“jrvwmp” – “microbe”) and طحالب (“THAlb” – “algae”); *inanimate objects* such as جزمة (“jzmp” – “shoes”) and سطل (“sTI” – “bucket”) among other usages.

**Simile and metaphor:** Users use simile and




Figure 3: Gender distribution for offensive language and its sub-categories




Figure 4: Distribution of Hate Speech Types. Note: A tweet may have more than one type.



Figure 5: Tag cloud for words with top valence score among offensive class, e.g. name calling (animals), curses, insults, etc.

metaphor were they would compare a person to: an *animal* as in زي الثور (“zy Alwvr” – “like a bull”), سمعني نهيك (“smEny nhyqk” – “let me hear your braying”), and هز ديلك (“hz dylk” – “wag your tail”); a person with *mental or physical disability* such as منغولي (“mngwly” – “Mongolian (Down syndrome)”), معوق (“mEwq” – “disabled”), and قزم (“qzm” – “dwarf”); and to the *opposite gender* such as جيش نوال (“jy\$ nwAl” – “Nawal’s army

(Nawal is female name)”) and نادي زيزي (“nAdy zyzy” – “Zizi’s club (Zizi is a female nickname)”).

**Indirect speech:** This includes: *sarcasm* such as إخوانك أذكى (“A\*kY AxwAtk” – “smartest one of your siblings”) and فيلسوف الحمير (“fylswf AIHmyr” – “the donkeys’ philosopher”); *questions* such as ايه كل الغباء ده (“Ayh kl AlgbA dh” – “what is all this stupidity”); and *indirect speech* such as النقاش مع البهايم غير مثمر (“AlnqA\$ mE AlbhAym gyr mvmr” – “no use arguing with cattle”).

**Wishing Evil:** This entails wishing death or major harm to befall someone such as ربنا ياخذك (“rbnA yAxdk” – “May God take (kill) you”), الله يلعنك (“Allh ylenk” – “may God curse you”), and روح في داهية (“rwH fy dAhyp” – equivalent to “go to hell”).

**Name alteration:** One common way to insult others is to change a letter or two in their names to produce new offensive words that rhyme with the original names. Some such examples include changing الجزيرة (“Aljzyrp” – “Aljazeera (channel)”) to الخنزيرة (“Alxnzyrp” – “the pig”) and خلفان (“xl-fAn” – “Khalfan (person name)”) to خرفان (“xrfAn” – “crazed”).

**Societal stratification:** Some insults are associated with: certain *jobs* such as بواب (“bwAb” – “doorman”) or خادم (“xAdm” – “servant”); and specific *societal components* such as بدوي (“bdwy” – “bedouin”) and فلاح (“flAH” – “farmer”).

**Immoral behavior:** These insults are associated with negative moral traits or behaviors such as حقير (“Hqyr” – “vile”), خاين (“xAyn” – “traitor”), and منافق (“mnAfq” – “hypocrite”).

**Sexually related:** They include expressions such as حول (“xwl” – “gay”), وسخة (“wsxp” – “prostitute”), and عرس (“ErS” – “pimp”).

Figure 5 shows the top words with the highest valence scores for individual words in the offensive tweets. Larger fonts are used to highlight words with highest scores and align as well with the categories mentioned in the breakdown for the offensive languages. We slightly modified the valence score described by (Conover et al., 2011) to magnify its value by multiplying valence with frequency of occurrence.

## 4 Experiments

We conducted an extensive battery of experiments on the dataset to establish strong Arabic offensive language classification results. Though offensive tweets have finer-grained labels where offensive tweet could also be vulgar and/or hate speech, we conducted coarser-grained classification to determine if a tweet was offensive or not. For classification, we experimented with several tweet representation and classification models. For tweet representations, we used: the count of positive and negative terms, based on a polarity lexicon; static embeddings, namely fastText and SkipGram; and deep contextual embeddings, namely BERT<sub>base-multilingual</sub> and AraBERT (Antoun et al., 2020).

### 4.1 Data Pre-processing

We performed several text pre-processing steps. First, we tokenized the text using the Farasa Arabic NLP toolkit (Abdelali et al., 2016). Second, we removed URLs, numbers, and all tweet specific tokens, namely mentions, retweets, and hashtags as they are not part of the language semantic structure, and therefore, not usable in pre-trained embeddings. Third, we performed basic Arabic letter normalization, namely variants of the letter *alef* to *bare alef*, *ta marbouta* to *ha*, and *alef maqsoura* to *ya*. We also separated words that are commonly incorrectly attached such as ياكلب (“yAkIb” – “O dog”), is split to يا كلب (“yA kIb”). Lastly, we normalized letter repetitions to allow for a maximum of 2 repeated letters. For example, the token ههههه (“hhhhh” – “hahahahaha”) is normalized to هه (“hh”). We also removed Arabic diacritics and word elongations (kashida).

### 4.2 Representations

**Lexical Features** Since offensive words typically have a negative polarity, we wanted to test the effectiveness of using a polarity lexicon in detecting offensive tweets. For the lexicon, we used NileULex (El-Beltagy, 2016), which is an Arabic polarity lexicon containing 3,279 MSA and 2,674 Egyptian terms, out of which 4,256 are negative and 1,697 are positive. We used the counts of terms with positive polarity and terms with negative polarity in tweets as features.

**Static Embeddings** We experimented with various static embeddings that were pre-trained on

different corpora with different vector dimensionality. We compared pre-trained embeddings to embeddings that were trained on our dataset. For pre-trained embeddings, we used: **fastText** Egyptian Arabic pre-trained embeddings (Bojanowski et al., 2017) with vector dimensionality of 300; **Ar-aVec** skip-gram embeddings (Mohammad et al., 2017), trained on 66.9M Arabic tweets with 100-dimensional vectors; and **Mazajak** skip-gram embeddings (Abu Farha and Magdy, 2019), trained on 250M Arabic tweets with 300-dimensional vectors. Sentence embeddings were calculated by taking the mean of the embeddings of their tokens. The importance of testing a character level n-gram model like fastText lies in the agglutinative nature of the Arabic language. We trained a new fastText text classification model (Joulin et al., 2017) on our dataset with vectors of 40 dimensions, 0.5 learning rate, 2–10 character n-grams as features, for 30 epochs. These hyper-parameters were tuned using a 5-fold cross-validated grid-search.

**Deep Contextualized Embeddings** We also experimented with pre-trained contextualized embeddings with fine-tuning for down-stream tasks. Recently, deep contextualized language models such as BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2019), UMLFIT (Howard and Ruder, 2018), and OpenAI GPT (Radford et al., 2018), have achieved ground-breaking results in many NLP classification and language understanding tasks. In this paper, we fine-tuned BERT<sub>base-multilingual</sub> (or simply BERT) and AraBERT embeddings to classify Arabic offensive language on Twitter as it eliminates the need for feature engineering. Although Robustly Optimized BERT (RoBERTa) embeddings perform better than (BERT<sub>large</sub>) on GLUE (Wang et al., 2018), RACE (Lai et al., 2017), and SQuAD (Rajpurkar et al., 2016) tasks, pre-trained multilingual RoBERTa models are not available. BERT is pre-trained on Wikipedia text from 104 languages, and AraBERT is trained on a large Arabic news corpus containing 8.5M articles composed of roughly 2.5B tokens. Both use identical architectures and come with hundreds of millions of parameters. Both contain an encoder with 12 Transformer blocks, hidden size of 768, and 12 self-attention heads. These embeddings use BP sub-word segments. Following Devlin et al. (2019), the classification consists of introducing a dense layer over the final hidden state  $h$  corresponding to first token

of the sequence, [CLS], adding a softmax activation on the top of BERT to predict the probability of the  $l$  label:  $p(l|h) = \text{softmax}(Wh)$ , where  $W$  is the task-specific weight matrix. During fine-tuning, all BERT/AraBERT parameters together with  $W$  are optimized end-to-end to maximize the log-probability of the correct labels.

### 4.3 Classification Models

We explored different classifiers. When using lexical features and pre-trained static embeddings, we primarily used an SVM classifier with a radial basis function kernel. Only when using the Mazajak embeddings, we experimented with other classifiers such as AdaBoost and Logistic regression. The SVM classifier performed the best on static embeddings, and we picked the Mazajak embeddings because they yielded the best results among all static embeddings. We used the Scikit Learn implementations of all the classifiers such as libsvm for the SVM classifier. We also experimented with fastText, which trained embeddings on our data. When using contextualized embeddings, we fine-tuned BERT and AraBERT by adding a fully-connected dense layer followed by a softmax classifier, minimizing the binary cross-entropy loss function for the training data. For all experiments, we used the PyTorch<sup>2</sup> implementation by HuggingFace<sup>3</sup> as it provides pre-trained weights and vocabularies.

### 4.4 Evaluation

For all of our experiments, we used 5-fold cross validation with identical folds for all experiments. Table 2 reports on the results of using lexical features, static pre-trained embeddings with an SVM classifier, embeddings trained on our data with fastText classifier, and BERT and AraBERT over a dense layer with softmax activation. As the results show, using fine-tuned AraBERT yielded the best results overall, followed closely by Mazajak/SVM, with large improvements in precision over using BERT. The success of AraBERT was surprising given that it was not trained on social media text. Perhaps, pre-training a Transformer model on social media text may improve results further. We suspect that the Mazajak/SVM combination performed better than BERT due to the fact that the Mazajak embeddings, though static, were trained

<sup>2</sup><https://pytorch.org/>

<sup>3</sup><https://github.com/huggingface/transformers>

on in-domain data, as opposed to BERT. For completeness, we compared 7 other classifiers with SVM using Mazajak embeddings. As results in Table 3 show, using SVM yielded the best results.

Model/classifier	Prec.	Recall	F1
Lexical Features			
SVM	68.5	35.3	46.6
Pre-trained static embeddings			
fastText/SVM	76.7	43.5	55.5
AraVec/SVM	85.5	69.2	76.4
Mazajak/SVM	<b>88.6</b>	72.4	79.7
Embeddings trained on our data			
fastText/fastText	82.1	68.1	74.4
Contextualized embeddings			
BERT <sub>base-multilingual</sub>	78.3	74.0	76.0
AraBERT	84.6	<b>82.4</b>	<b>83.2</b>

Table 2: Classification performance with different features and models.

Model	Prec.	Recall	F1
Decision Tree	51.2	53.8	52.4
Random Forest	82.4	42.4	56.0
Gaussian NB	44.9	<b>86.0</b>	59.0
Perceptron	75.6	67.7	66.8
AdaBoost	74.3	67.0	70.4
Gradient Boosting	84.2	63.0	72.1
Logistic Regression	84.7	69.5	76.3
SVM	<b>88.6</b>	72.4	<b>79.7</b>

Table 3: Performance of different classification models on Mazajak embeddings.

#### 4.5 Error Analysis

We inspected the tweets of one fold that were misclassified by the Mazajak/SVM model (36 false positives/121 false negatives) to determine the most common errors. They were as follows:

##### Four false positive types:

- **Gloating:** ex. يا هييده (“yA hbydp” - “O you delusional”) referring to fans of rival sports team for thinking they could win.
- **Quoting:** ex. لما حد يسب ويقول يا كلب (“lma Hd ysb wyqwl yA klb” - “when someone swears and says: O dog”).
- **Idioms:** ex. يا فاطر رمضان يا خاسر دينك (“yA fATr rmDAn yA xAsr dynk” - “o you who does not fast Ramadan, you have lost your faith”), which is a colloquial idiom.

- **Implicit Sarcasm:** ex. يا خاين انت عايز تشكك في حب الشعب للرئيس 😂😂 (“yA xAyn Ant EAww t\$kk fy Hb Al\$Eb llrys” - “O traitor, (you) want to question people’s love for the president 😂😂”) where the author is mocking the president’s popularity.

##### Two false negative types:

- **Mixture of offensiveness and admiration:** ex. calling a girl a puppy يا كلبوة (“yA klbwbp” - “O puppy”) in a flirtatious manner.
- **Implicit offensiveness:** ex. calling for cure while implying sanity: وتشفي حكام بلدك من المرض (“wt\$fy HkAm bldk mn AlmrD” - “and cure rulers of your country from illness”).

## 5 Conclusion and Future Work

In this paper we presented a systematic method for building an Arabic offensive language tweet dataset that does not favor specific dialects, topics, or genres. We developed detailed guidelines for tagging the tweets as clean or offensive, including special tags for vulgar tweets and hate speech. We tagged 10,000 tweets, which we plan to release publicly and would constitute the largest available Arabic offensive language dataset. We characterized the offensive tweets in the dataset to determine the topics that illicit such language, the dialects that are most often used, the common modes of offensiveness, and the gender distribution of their authors. We performed this breakdown for offensive tweets in general and for vulgar and hate speech tweets separately. We believe that this is the first detailed analysis of its kind. Lastly, we conducted a large battery of experiments on the dataset, using cross-validation, to establish a strong system for Arabic offensive language detection. We showed that using an Arabic specific BERT model (AraBERT) and static embeddings trained on tweets produced competitive results on the dataset.

For future work, we plan to pursue several directions. First, we want explore target specific offensive language, where attacks against an entity or a group may employ certain expressions that are only offensive within the context of that target and completely innocuous otherwise. Second, we plan to examine the effectiveness of cross dialectal and cross lingual learning of offensive language.



## References

- Ahmed Abdelali, Kareem Darwish, Nadir Durrani, and Hamdy Mubarak. 2016. Farasa: A fast and furious segmenter for arabic. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Demonstrations*, pages 11–16.
- Ehab Abozinadah. 2017. *Detecting Abusive Arabic Language Twitter Accounts Using a Multidimensional Analysis Model*. Ph.D. thesis, George Mason University.
- Ibrahim Abu Farha and Walid Magdy. 2019. **Mazajak: An online Arabic sentiment analyser**. In *Proceedings of the Fourth Arabic Natural Language Processing Workshop*, pages 192–198, Florence, Italy. Association for Computational Linguistics.
- Sweta Agrawal and Amit Awekar. 2018. Deep learning for detecting cyberbullying across multiple social media platforms. In *European Conference on Information Retrieval*, pages 141–153. Springer.
- Azalden Alakrot, Liam Murray, and Nikola S Nikolov. 2018. Towards accurate detection of offensive language in online communication in arabic. *Procedia computer science*, 142:315–320.
- Nuha Albadi, Maram Kurdi, and Shivakant Mishra. 2018. Are they our brothers? analysis and detection of religious hate speech in the arabic twitter-sphere. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 69–76. IEEE.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. Arabert: Transformer-based model for arabic language understanding. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15.
- Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. 2017. Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 759–760. International World Wide Web Conferences Steering Committee.
- Pablo Barberá and Gaurav Sood. 2015. Follow your ideology: Measuring media ideology on social networks. In *Annual Meeting of the European Political Science Association, Vienna, Austria*. Retrieved from <http://www.gsood.com/research/papers/mediabias.pdf>.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Thomas Chadeaux. 2014. Early warning signals for war in the news. *Journal of Peace Research*, 51(1):5–18.
- Michael Conover, Jacob Ratkiewicz, Matthew R Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. *ICWSM*, 133:89–96.
- Kareem Darwish, Dimitar Alexandrov, Preslav Nakov, and Yelena Mejova. 2017. Seminar users in the arabic twitter sphere. In *International Conference on Social Informatics*, pages 91–108. Springer.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Eleventh International Conference on Web and Social Media (ICWSM)*, pages 512–515.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. 2015. Hate speech detection with comment embeddings. In *Proceedings of the 24th international conference on world wide web*, pages 29–30. ACM.
- Samhaa R. El-Beltagy. 2016. **NileULex: A phrase and word level sentiment lexicon for Egyptian and modern standard Arabic**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 2900–2905, Portorož, Slovenia. European Language Resources Association (ELRA).
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Jeremy Howard and Sebastian Ruder. 2018. **Universal language model fine-tuning for text classification**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Melbourne, Australia. Association for Computational Linguistics.
- Timothy Jay and Kristin Janschewitz. 2008. The pragmatics of swearing. *Journal of Politeness Research. Language, Behaviour, Culture*, 4(2):267–288.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431. Association for Computational Linguistics.
- Irene Kwok and Yuzhou Wang. 2013. Locate the hate: Detecting tweets against blacks. In *Twenty-seventh AAAI conference on artificial intelligence*.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. **Race: Large-scale reading comprehension dataset from examinations**.
- Alessandro Maisto, Serena Pelosi, Simonetta Vietri, Pierluigi Vitale, and Via Giovanni Paolo II. 2017. Mining offensive language on social media. *CLiC-it 2017 11-12 December 2017, Rome*, page 252.

- Shervin Malmasi and Marcos Zampieri. 2017. Detecting hate speech in social media. *arXiv preprint arXiv:1712.06427*.
- Elisa Mattiello. 2005. The pervasiveness of slang in standard and non-standard english. *Mots Palabras Words*, 5:7–41.
- Abu Bakr Mohammad, Kareem Eissa, and Samhaa El-Beltagy. 2017. *Aravec: A set of arabic word embedding models for use in arabic nlp*. *Procedia Computer Science*, 117:256–265.
- Hamdy Mubarak and Kareem Darwish. 2019. Arabic offensive language classification on twitter. In *International Conference on Social Informatics*, pages 269–276. Springer.
- 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.
- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pages 145–153. International World Wide Web Conferences Steering Committee.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *URL [https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language\\_understanding\\_paper.pdf](https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf)*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. *Squad: 100,000+ questions for machine comprehension of text*.
- Younes Samih, Mohamed Eldesouki, Mohammed Attia, Kareem Darwish, Ahmed Abdelali, Hamdy Mubarak, and Laura Kallmeyer. 2017. Learning from relatives: unified dialectal arabic segmentation. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 432–441.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. *Glue: A multi-task benchmark and analysis platform for natural language understanding*.
- Zeeraq Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of the NAACL student research workshop*, pages 88–93.
- Dawei Yin, Zhenzhen Xue, Liangjie Hong, Brian D Davison, April Kontostathis, and Lynne Edwards. 2009. Detection of harassment on web 2.0. *Proceedings of the Content Analysis in the WEB*, 2:1–7.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). *arXiv preprint arXiv:1903.08983*.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. Semeval-2020 task 12: Multilingual offensive language identification in social media (offenseval 2020). *arXiv preprint arXiv:2006.07235*.