

ARTICLE

# Emojis as anchors to detect Arabic offensive language and hate speech

Hamdy Mubarak<sup>1</sup>, Sabit Hassan<sup>2</sup> and Shammur Absar Chowdhury<sup>1</sup> 

<sup>1</sup>Qatar Computing Research Institute, Hamad Bin Khalifa University, Doha, Qatar and <sup>2</sup>School of Computing and Information, University of Pittsburgh, Pittsburgh, PA, USA

**Corresponding author:** Hamdy Mubarak; Email: [hmubarak@hbku.edu.qa](mailto:hmubarak@hbku.edu.qa)

(Received 1 September 2021; revised 31 March 2022; accepted 19 May 2022; first published online 10 August 2023)

## Abstract

We introduce a generic, language-independent method to collect a large percentage of offensive and hate tweets regardless of their topics or genres. We harness the extralinguistic information embedded in the emojis to collect a large number of offensive tweets. We apply the proposed method on Arabic tweets and compare it with English tweets—analyzing key cultural differences. We observed a constant usage of these emojis to represent offensiveness throughout different timespans on Twitter. We manually annotate and publicly release the largest Arabic dataset for *offensive*, *fine-grained hate speech*, *vulgar*, and *violence* content. Furthermore, we benchmark the dataset for detecting offensiveness and hate speech using different transformer architectures and perform in-depth linguistic analysis. We evaluate our models on external datasets—a Twitter dataset collected using a completely different method, and a multi-platform dataset containing comments from Twitter, YouTube, and Facebook, for assessing generalization capability. Competitive results on these datasets suggest that the data collected using our method capture universal characteristics of offensive language. Our findings also highlight the common words used in offensive communications, common targets for hate speech, specific patterns in violence tweets, and pinpoint common classification errors that can be attributed to limitations of NLP models. We observe that even state-of-the-art transformer models may fail to take into account culture, background, and context or understand nuances present in real-world data such as sarcasm.

**Keywords:** Offensive language; Hate speech; Emojis; Text classification; Social media analysis

## 1. Introduction

**Disclaimer:** *Due to the nature of this work, some examples contain offensiveness, hate speech, and profanity. This doesn't reflect the authors' opinions by any means. We hope this work can help in detecting and preventing the spread of such harmful content.*

Social media platforms provide a medium for individuals or a group to connect with the world and share their opinion (Intapong *et al.* 2017), often with limited inhibitions. Taking the advantage of such liberation, social media users may use vulgar, pornographic, or hateful language. Such behavior can result in the spread of verbal hostility and can impact users' psychological well-being (Gülaçtı 2010; Waldron 2012).

In recent years, Twitter has become highly popular in the Arab region (Abdelali *et al.* 2021b), with more than 27 million tweets per day (Alshehri *et al.* 2018). Many look to Twitter to express their views, ideas, and share their stories. While the importance of sharing these views and ideas is immense, the aforementioned societal problem of sharing malicious content also arises.



With the sheer volume of content on social media, manually filtering out malicious content while maintaining users' right to freedom of expression is virtually impossible for the platform providers. The increased risk and effect of such hostility presence in social media has attracted many multidisciplinary researchers and motivated the need to automatically detect offensiveness of posts/comments and utilize the system to: (i) filter adult content (Cheng *et al.* 2014; Mubarak *et al.* 2021); (ii) quantify the intensity of polarization (Belcastro *et al.* 2020; Conover *et al.* 2011); (iii) classify trolls and propaganda accounts (Alhazbi 2020; Darwish *et al.* 2017; Dimitrov *et al.* 2021a, b); and (iv) identify hate speech and conflicts (Kiela *et al.* 2020; Davidson *et al.* 2017; Ousidhoum *et al.* 2019; Chung *et al.* 2019).

Machine learning-based automation approaches require labeled datasets to distinguish malicious content from others. Constructing datasets, however, is not a straightforward process. In a randomly sampled collection of tweets, the percentage of malicious tweets is very small. Mubarak *et al.* (2017) report that only 1–2% of Arabic tweets are abusive. The percentage of hate speech is even smaller. This implies that to obtain a sizable dataset of malicious content, a massive number of tweets need to be labeled. To avoid this, existing approaches utilize heuristics to increase the percentage of malicious tweets before manually labeling them. Such heuristics include searching for specific language-dependent keywords or patterns.

In this paper, we present an automated emoji-based approach to collecting tweets that have a much higher percentage of malicious content, without having any language dependency. Emojis have quickly become an important part of our daily communication. They are used worldwide for conveying messages without any language or online platform barriers and have been referred to as a universal language (Mei 2019; Dürscheid and Siever 2017). Hence, the extralinguistic information carried by the emojis can be instrumental in capturing offensive content, without being impacted by the lack of language-dependent knowledge, or any preferred linguistic patterns. Using emojis also resolves challenges posed by nonstandard spellings of offensive words.

To this end, we use emojis as anchors to collect tweets and manually annotate them for fine-grained abusive/offensive language categories—including hate speech, vulgar, and violent tweet content. We exploit the collected dataset for studying: (i) emoji usage in a different period of time; (ii) extraction of offensive words with different dialectal spellings and morphological variations; (iii) hate speech targets; (iv) linguistic content in vulgar and profanity words; and (v) exploring common patterns present in violent tweets.

Our comparative analysis shows that by using emojis as anchors, we end up with 35% offensive (OFF) tweets and 11% hate speech (HS) tweets, which is almost as double the approach used by Mubarak *et al.* (2017). Our experiments with various machine learning and deep learning classifiers show that the classifiers are able to learn effectively from the data. Further, we show that these classifiers are capable of generalizing well on two external datasets: (i) the OffensEval 2020 (SemEval, task 12) dataset (Zampieri *et al.* 2020) for Arabic offensive language identification, and (ii) the MPOLD dataset (Chowdhury *et al.* 2020b)—a multi-platform dataset containing news comments annotated for offensiveness from Twitter, YouTube, and Facebook. This suggests that our method captures universal characteristics of offensive language on social media and can be used to collect larger datasets with less effort and linguistic knowledge support.

Therefore, the main contributions of the paper are as follows:

- We present an emoji-based method to collect offensive and hate speech tweets. We show that using the method, we can collect a higher percentage of offensive and hate speech content compared to the existing methods. Moreover, we demonstrate that using this method we can collect a large number of offensive content in other languages such as Bengali.
- We manually label the largest dataset for offensiveness, along with fine-grained hate speech types, vulgarity (profanity), and violence. This labeled data will be publicly released for the community.<sup>a</sup>

<sup>a</sup>Data can be downloaded from this URL: <https://sites.google.com/view/arabichate2022>.

- We perform an in-depth analysis of different properties of the dataset including common offensive words and hate speech targets. We show also that there are some cultural differences in using emojis in Arabic offensive tweets compared to English.
- We build effective machine learning and deep learning-based classifiers to automatically tag offensive and hate speech with high accuracy. We show that our model perform well on external datasets collected from Twitter and also demonstrate the potential generalization capability of the designed model in different social media platforms—such as YouTube and Facebook.
- We analyze common classification errors and provide some recommendations to enhance model explainability.

While our main focus, in this paper, is to create an Arabic offensive dataset, this method can easily be adapted for other languages and potentially other text classification tasks such as sentiment analysis or emotion detection.

The paper is structured as follows. In section 2, we discuss related work, focusing on the methods used by other researchers to collect datasets for offensive language. In section 3, we describe our data collection and annotation jobs. Section 4 contains an extensive analysis of the dataset. The analysis includes a study of offensive emojis, and how emoji usage changes in a different time period. We further analyze offensive words and hate speech targets found in our dataset. In Section 5, we train a set of machine learning and deep learning models for the classification of offensive language and hate speech. We manually analyze errors made by our models and also use Local Interpretable Model-Agnostic Explanations (LIME) explainability tool to interpret decisions made by our models. In this section, we also show that models trained on our data achieve good results on other datasets. Section 6 discusses ethics and the social impact of our work. Finally, in Section 7, we present a summary of our findings.

## 2. Related work

There has been a large amount of research in recent years to address and detect the growing use of “offensive languages” and hate speech in different social platforms (see Nakov *et al.* (2021) and Salminen *et al.* (2020) for more details). Detecting offensive language has been the focus of many shared tasks such as OffensEval 2020 (Zampieri *et al.* 2020) for five languages and OSACT 2020 (Mubarak *et al.* 2020b) for Arabic. The best systems at the shared tasks (Alami *et al.* 2020; Hassan *et al.* 2020) utilized support vector machines (SVM) and fine-tuned transformer models. Hate speech has been less explored compared to offensive language. A dataset consisting of 5% hate speech was presented at OSACT 2020 shared task. The best system performed extensive preprocessing including normalizing **emojis** (translate their English description to Arabic) and dialectal Arabic (DA) to modern standard Arabic (MSA) conversion among others (Husain 2020). ASAD (Hassan *et al.* 2021a) is an online tool that utilizes the shared task datasets for offensiveness and hate speech detection in tweets along with other social media analysis components such as emotion (Hassan *et al.* 2021b) and spam detection (Mubarak *et al.* 2020a).

To collect potentially offensive tweets, some studies use a list of seed offensive words and hashtags (e.g., Mubarak *et al.* 2017). This approach has several drawbacks: (i) it is hard to maintain this list as offensive words are ever-evolving; (ii) Arabic dialects are widely used on social media and building such a list for different dialects is a challenging task that requires deep knowledge about cultures in many countries; and (iii) Arabic dialects have no standard writings. So, it is extremely difficult to list all possible surface forms and creative spellings for words in general including offensive words. Moreover, (iv) Arabic has a rich morphology (both derivational and inflectional), and a large number of affixes can be attached to words—for example **وسيفعلها** (“wsyfElhA”—“and he will do it”)<sup>b</sup>—which makes string match for offensive words less optimal.

<sup>b</sup>We provide Arabic examples and their Buckwalter transliteration and English translation throughout this paper.

**Table 1.** Available Arabic datasets along with the annotation labels, data source and collection method, and percentage of offensive content.

Dataset	Source	Method	Size	Language	OFF%	Annotations (values (1: Yes, 0: No))
Mubarak <i>et al.</i> (2017)	TW	Controversial users	1100	Egyptian	59%	OFF (1/0), Vulgar (1/0)
	AJ	Deleted comments	31,692	MSA	82%	OFF (1/0), Vulgar (1/0)
Albadi <i>et al.</i> (2018)	TW	Keywords	6030	MSA/DA	45%	Religious HS (1/0)
Alshaalan and Al-Khalifa (2020)	TW	Keywords	9316	Saudi	28%	HS (1/0)
Mubarak <i>et al.</i> (2020b)	TW	Pattern	10,000	MSA/DA	19%	OFF (1/0), HS (1/0)
Chowdhury <i>et al.</i> (2020b)	TW, FB, YT	User replies to AJ	4000	MSA/DA	17%	OFF (1/0), HS (1/0), Vulgar (1/0)
Our method	TW	Emojis	12,698	MSA/DA	35%	OFF (1/0), Vulgar (1/0), Violence (1/0) HS (gender, race, ideology, social class, religion, disability)

Sources: TW (Twitter), Aljazeera (AJ), FB (Facebook), YT (YouTube).

Finally, (v) the offensiveness of many words is highly dependent on context, for example, the word كلب (“klb”—“dog”) can be used in offensive contexts such as هو كلب (“hw klb”—“He is a dog”) and in clean contexts such as عندي كلب (“Endy klb”—“I have a dog”).

Authors in Mubarak *et al.* (2020c) showed that offensive language exists in less than 2% of any random collection of tweets, and by considering a common pattern used in offensive communications, this ratio increases to 19% (5% are hate speech). This pattern is يا .. يا (“yA .. yA”—“O .. O ..”) which is used mainly to direct the speech to a person or a group. This pattern is used across all dialects without any preference for topics or genres; however, it cannot be generalized to other languages.

In the same research, authors observed that the most frequent personal attack on Arabic Twitter is to call a person an animal name (i.e., direct name-calling), and the most used animals are كلب (“klb”—“dog”), حمار (“HmAr”—“donkey”), and بهيم (“bhym”—“beast”) among others. In Chowdhury *et al.* (2020b), authors analyzed Arabic offensive language on Facebook, YouTube, and Twitter, and they showed that some emojis are widely used in offensive communications, including some animals (dog, pig, monkey, cow, etc.), some face emojis (anger, disgust, etc.), and others (shoe, etc.).

Unlike the aforementioned studies, we propose a generic method to collect offensive tweets regardless of their topics, genres, or dialects. We applied it to Arabic Twitter and obtained the largest dataset of tweets labeled for offensiveness (approx. 13K tweets) with 35% of the tweets labeled as offensive and 11% as hate speech. In this method, we simply use a list of emojis that commonly appear in offensive communications, including emojis that express anger, some animals, and inanimate things extracted from existing datasets of offensive tweets. We believe with small modifications to this list, considering cultural differences, the approach can be used for other languages as well to collect a large number of offensive tweets. Comparison with available Arabic datasets for offensive language and hate speech is shown in Table 1.

It is worth noting that authors in Wiegand and Ruppenhofer (2021) used emojis to detect English abusive tweets. They used eight emojis based on correlations between concepts and



Annotators fully agree (3 out of 3) in 65% of the tweets. For better quality, we hired and trained an expert linguist who is familiar with different dialects to carefully check all tweets having different opinions. This resulted in changing 18% of the labels (w.r.t to the previously assigned majority label) for the tweets.

### 3.3 Annotation of hate speech, profanity, and violence

We created another job on Appen with all the annotated offensive tweets from *Job1* and asked annotators to detect the presence of hate speech. We defined hate speech as any content that contains offensive language targeting a group of people based on common characteristics, such as race/ethnicity/nationality, religion/belief, ideology, disability/disease, social class, and gender. We used job settings similar to *Job1*. Additionally, we asked annotators to judge whether the offensive tweets have profanity or vulgar words and whether they promote violence. Basic statistics and examples from different annotations are listed in Table 2.

## 4. Analysis

We obtained a total of 582 different emojis in the final annotated data. Out of those, 122 emojis appear more than 20 times. In this section, we first report some observations seen in our data, followed by in-depth analysis of emoji in a different time period, linguistic characteristics associated with emojis along with targets of hate speech in these tweets.

### 4.1 Common offensive emojis

We report our observations about common emojis used to represent offensiveness below:

- There are additional emojis, unique to our initial seed list, frequently used in the annotated offensive tweets. This class of emojis (see “New emojis” row in Figure 3) includes instances like spit/drops (65% of tweets having it was labeled as offensive) to hammer (18% offensive). This suggests a potential for enriching the initial emoji seed list and expanding the dataset in an iterative way.
- Top offensive emojis are shown in Figure 3 (mostly animals). If a tweet has any of them, most likely it is labeled as offensive [pig (84% offensive) to middle finger (62%)]. Similar emojis appear in hate speech tweets with percentages range from 36% to 18%, respectively.
- Top vulgar emojis are mostly used in tweets having adult content (e.g., adult ads).
- For violence category, the most common emojis are knife, punch, and shoe in order.

### 4.2 Emojis usage in a different time period

We collected tweets from a different time period (in March 2021) to study the usage pattern of offensive emojis over time. For this, we selected six emojis from Figure 3 and their corresponding words to extract their tweets.<sup>e</sup> For each emoji and its corresponding word, we then randomly selected 200 tweets and asked an Arabic native speaker to judge the content for offensiveness.<sup>f</sup>

Our study suggests that some emojis (see Figure 4) and their corresponding words are widely used in offensive tweets during different periods of time. While the middle finger was reported as

<sup>e</sup>We ignored “high heel shoe” emoji as it is used dominantly in adult ads, and “donkey/horse” emoji as it is ambiguous in its usage (as donkey in some cases and horse in others). “Donkey” emoji is not found in <https://emojipedia.org>

<sup>f</sup>Some words appear less than 200 times in this collection.

**Table 2.** Statistics and examples from the annotated corpus (total of 12,698 tweets)

Class/subclass	#	%	Example
<b>Clean</b> (CLN)	8,235	65	لن تحصل على غدٍ أفضل مادمت تفكر بالأمس You won't have a better tomorrow as long as you think about yesterday.
<b>Offensive</b> (OFF)	4,463	35	يلعن ابوك على هالسؤال. عساه ينقرض الكريه May God curse your father for this question! I hope this fool will die out!
* Hate speech	1,339	11	(Note: 30% of offensive tweets are labeled as hate speech)
- Gender	641	48	بنات اليوم قليلات أدب. والله ما نوصل لعمر بعض الرجال Girls today are impolite. I swear to God, we don't reach for some men immorality.
- Race	366	27	شعبكم متخلف. الله ياخذك إنتي والفلبين People of your country are musty. May God take (kill) you and the Philippines.
- Ideology	190	14	ناديك وضع لا شك في ذلك. حزبك لا يقدر إلا على النباح Your club is vile, no doubt about that. Your party cannot do anything except barking.
- Social class	101	8	دامك مقيم انكتم وخل اهل الأرض الاصليين يتكلمون. ابلع يا سباك! As you are a resident, shut up and let original citizens speak. Swallow, plumber!
- Religion	38	3	إنتوا بتعملوا ف ديك أبونا كده ليه هو إحنا كفره ولا يهود Why are you doing this to us? Are we disbelievers or Jews?
- Disability	3	0	ذا القزم طلعت جايزتن له بس ماعرف يعبر This dwarf got two prizes, but he does not know how to express.
* Vulgar	189	1.5	كاتبة ع حسابك قبة بنت قبة و بتعجبي انهم يبيعشو صور فاضحة * قبة: امرأة فاجرة فأيدة تمارس الإغواء (معجم اللغة العربية المعاصر) You describe yourself as a prostitute, and you wonder why they're sending porn photos!
* Violence	85	0.7	صفعه بهذا الحذاء حجم راسك على قرعتك. يبيجع ودي اطعنها Slap by this shoe on your bald head. Yucky! I'd like to stab her.

the strongest offensive emoji in English tweets (Wiegand and Ruppenhofer 2021), we found that there are more common offensive emojis that appear in our dataset of Arabic tweets such as pig (84%), shoe (77%), and dog (68%) (first in “Top offensive emojis” row in Figure 4). Tweets having middle finger emoji were tagged as offensive in 62% of the cases. We found annotation errors due to short context in some cases, but some users use the middle finger mistakenly instead of index finger<sup>8</sup> as shown in Figure 5.

### 4.3 Emojis and linguistic usage

As observed in Donato and Paggio (2017), emojis are often redundant and convey something already expressed verbally in tweets. Thus, offensive emojis may co-occur with many

<sup>8</sup>We anticipate middle finger emoji is not widely understood as offensive by all Arab users.

New emojis	
Top offensive emojis	
Top hate speech emojis	
Top vulgar emojis	
Top violence emojis	

Figure 3. Categories of common offensive Emojis. Emojis in the same row are sorted based on percentage of offensive tweets (in descending order).

Emoji	#	OFF%	Word	#	OFF%
	200	72	خنزير pig	8	75
	200	86	جزمة shoe	200	86
	200	84	كلب dog	200	91
	200	77	خروف sheep	200	80
	200	37	تفلة/بصقة spit	169	98
	200	64	بقرة cow	200	75

Figure 4. Percentage of offensive tweets for emojis and their corresponding words in samples from March 2021.

 حامد تركي بويابس @HBuyabes · Feb 25, 2018  
 ((نَمْ لِنَسْأَلَنَّ يَوْمِيذٍ عَنِ النَّعِيمِ))  
 صدق الله العظيم  
 الصورة 🍌  
 أطفال الكويت  
 الصورة 🍌  
 أطفال أفريقيا  
 يا رب..  
 أسالك المغفرة والرحمة بنا

USER@ حاطة 🍌 بدل 🍌 🍌 🍌 reply to another tweet  
 Translation: You put 🍌 instead of 🍌 🍌 🍌

Figure 5. Wrong usage of middle finger emoji in Arabic tweets.

Dialect variation	يا كلب 🐕. تدري انه جلب 🐕 (yA klb. tdry Anh JIb) O dog 🐕! You know he is a dog 🐕
Number (Plural)	ياشوية 🐕 🐕 O bunch of 🐕 🐕! (dogs)
Gender (Feminine)	🐕ة وماتت She is a dog (f.) and died

Figure 6. Emojis co-occur with dialectal and morphological variations.

offensive words. As DA is widely used on Twitter, it's very hard to list all variations of offensive dialectal words. Figure 6 shows an example where offensive emojis co-occurred with offensive dialectal words which confirms the aforementioned observation. It also shows how emojis are used to cover morphological variations such as number and gender.



Table 3. Common targets in hate speech tweets

Target	Words	Example
Women	المرءات، الحريم Women, girls	اللله يسامح اللله سمح للحريم يدخلون الملعب. منظر مقرف! May God forgive those who allowed women to enter the stadium. Disgust!
Iran	إيران، الفرس، الماوس Iran, Persians, Magi	كل زق أنت وإيران الماوسية Eat shit, you and the Magian Iran.
Israel	إسرائيل، الصهاينة Israel, Zionists	الللعنة عليك وعلى إسرائيل May God curse you and Israel!
Jews	اليهود Jews	أنتم تقتلون الأبرياء وتتركون اليهود يا خونة You kill innocent people and leave Jews. O traitors!
Houthi (group)	الحوثيين، الحوثي Houthi, Houthis	دمرتم كل جميل في وطني. الحوثي عدو الانسانية You destroyed all the beauties in my country. Houthi is the enemy of humanity.
Saudi Arabia	السعودية، السعوديين KSA, Saudis	السعودية من الاحتضار الديني الى حضارة العري Saudi Arabia from religious agony to a civilization of nudity
Muslim brotherhood (group, party)	الإخوان Brotherhood	هادو المرتقة من أتباع حماس والإخوان These mercenaries are followers of Hamas and the Brotherhood
Qatar	قطر، القطريين Qatar, Qataris	اعلام دولة قطر تنبح ليل نهار Media in the tiny country of Qatar barks day and night
USA	أمريكا، الأمريكان America, Americans	خرا عليك وعلى روسيا وأمريكا وخود الخليج بمعبتك Shit on you and on Russia and the US in addition to the Gulf.
Turkey	تركيا، الأتراك، الترك، العثمانيين Turkey, Turks, Ottomans	أنت يا تركي كذاب مثل أجدادك You, Turkish, are a liar like your ancestors.

offensive words. However, the remaining 29% are not necessarily offensive and can appear in many clean contexts. This includes named entities such as *إيران، بوتين، الحوثي، بشار* (“AyrAn, bwtyN, AlHwvy, b\$Ar”—Iran, Putin, Houthi, Bashar) in addition to some words that were concentrated by chance in our offensive tweets. Examples of such words are *ذکور، بلدك، ذل، \*l, bldk, \*kwr*—humiliation, your country, males).

#### 4.5 Hate speech targets

We used the aforementioned valence score formula to extract common words appear in all hate speech tweets. Then, we manually filtered and grouped them to extract common targets for hate speech. These targets include women, countries, groups, and political parties as listed in Table 3.

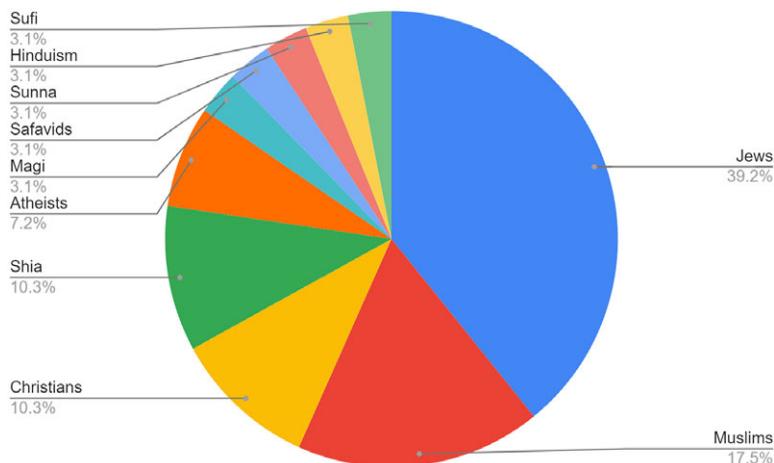


Figure 8. Hate speech% for different religious groups.

#### 4.5.1 Religious hate speech targets

Detecting religious hate speech in Arabic was the main focus of Albadi *et al.* (2018). Authors mentioned that Arabic is the official language in 6 of the 11 countries with the highest Social Hostilities Index.<sup>1</sup> They found that 33% of all hateful tweets targeted against Jews followed by Shia and Christians. We analyzed hateful tweets in our dataset, and we found similar results where Jews are the main target of religious prejudice with 39% of all hateful tweets followed by Muslims, Christians, and Shia. Figure 8 shows the distribution of hate speech tweets for different religious groups.

#### 4.6 Vulgar and profanity words

We observe that vulgar words include genitals, sexual actions, and references in addition to words used in adult ads. They are mainly written in DA, and the most common words are shown in Figure 7.

#### 4.7 Violence tweets

As violence tweets are rare (0.7% in our collection), we manually analyzed them all to extract common patterns. These patterns are summarized in Table 4. First, we define *jhead*<sub>z</sub> for head and neck, and *jbody*<sub>z</sub> for body parts including *jhead*<sub>z</sub> that can be used as objects for many violence verbs, for example, I will stab [him/his neck]. For future work, we plan to extend this list using synonyms and related words from dictionaries/BERT models. Then, we extract tweets having any of these verbs and targets after applying some morphological expansions.

#### 4.8 Is emojis a universal anchor?

The typical approach to crawl data from social media, especially from Twitter, is heavily dependent on task-specific keywords. Preparing such keyword-specific anchors needs significant knowledge of the language usage. However, using emojis to collect the data removes the initial language dependency and inherited bias from the keyword-search space.

<sup>1</sup>SHI is a measure of crimes partially motivated by religion.

**Table 4.** Common patterns in violence tweets

Pattern. S:Sbj, V:Verb, O:Obj	Description/verbs
<head>	رأس، دماغ، وجه، رقبة head, brain, face, neck
<body>	رأس، دماغ، وجه، رقبة، عين، خشم، أسنان، بطن head, brain, face, neck, eye, nose, teeth, stomach
SV O, O: <human> (ex: I will kill you)	يقتل، يذبح، يدمج kill, slaughter
SV O, O: <human>/<body> (ex: I will hit your head)	يدوس، يدعس، يجلد، يطعن، يضرب، يكسر trample, whip, stab, hit, break
SV O, O: <head> (ex: I will cut your neck)	يقطع، يفتح، يطير cut
<hit> on <body> (ex: A slap on your face)	كف، جزمة، صفعه: على shoe, slap

**Food:** You have salad and I'll have a big burger 🍔  
**Offensive:** 🐷 🐷 🐷. This man worst than shit.  


---

**Clothes:** So, get out your dancing shoes 👞 and get your tickets!  
**Offensive:** ...When wl u apply ur mind & recognize traitors 🖐️🖐️🖐️👞👞

**Figure 9.** Total percentage of offensive tweets in a sample of 50 tweets for each emoji in Bengali tweets.

In this section, we study the efficacy of emojis to collect offensive tweets for other languages. We utilize the proposed pipelines for Bengali (a low-resource language)<sup>j</sup> and analyzed the percentage of offensive tweets obtained using the method. Moreover, we also include analysis of English tweets to indicate how offensive communication and emoji usage are still culture-dependent if not language.

For Bengali, we collected 200 tweets in August 2021 for five emojis (a total of 1K tweets) from the different groups listed in Figure 2. From these 200, we randomly selected 50 tweets per emoji (a total of 250 tweets) and a Bengali native speaker annotated them for offensiveness. We noticed that ≈ 34% of all the annotated tweets are offensive. We observed that tweets with “middle finger” had the most frequent (50% of tweets) number of offensive content, followed by “pile of poo” emoji (32% of tweets). The detailed distribution is in Figure 9 with examples in Figure 10. We plan to use this method to collect and annotate offensive tweets for other low-resource languages such as Hindi,<sup>k</sup> in addition to Bengali. We keep this as a potential future work.

<sup>j</sup>Bengali is the lingua franca of Bangladesh and parts of India. There are 228 million native speakers, making it the fifth most-spoken native language in the world. Source: [https://en.wikipedia.org/wiki/Bengali\\_language](https://en.wikipedia.org/wiki/Bengali_language)

<sup>k</sup>Pilot annotation of Hindi tweets shows similar trends (52% offensive for “middle finger,” and 40% offensive for “angry face”).

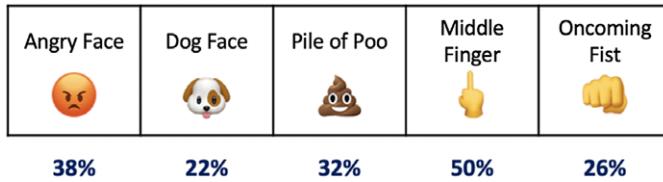


Figure 10. Examples of (non-)offensive tweets in Bengali collected using emojis.

**Bangla Non-Offensive Example:**

@xxxx ওরে যে আমার কি পরিমাণে বিরক্তিকর লাগে তোমারে আমি বইলা বুঝাইতে পারুম না 🤔🙄😞

**English Translation:**

@xxxx I can't even begin to explain to you how annoying I find him/her 🤔🙄😞

**Bangla Offensive Example:**

#স্বার্থপর হউল সুখে থাকুল 😄 Shorif Bro 🙄 যার চিন্তা সেই করে 😞 কাউরে #চোদার Tm নাই 🙄👉

**English Translation:**

be #selfish and be happy 😄 Shorif Bro 🙄 everyone thinks for themselves. 😞 cant be bothered to give a #f\*\*\* about anyone else. 🙄👉

Figure 11. English examples of pig and shoe emojis (topics of common usages and offensive contexts are shown).

As the usage of emojis for offensive content in English is well studied, we only collected English tweets having the top Arabic offensive emojis (pig and shoe) in April 2021. We collected more than 8,000 tweets per each, and we randomly selected and annotated 200 tweets for each emoji. We noticed that less than 5% of those tweets are offensive. Often, pig emoji appears in contexts about food while shoe emoji appears in tweets talking about fashion and clothes. Examples are shown in Figure 11.

This suggests that while some emojis are widely used in offensive communications in a particular culture, they might be used completely in different ways in other languages or cultures. Thus, it's important to keep these differences in mind when collecting data using emojis from different cultures.

**5. Experiments and results**

In this section, we describe the details of machine learning algorithms—SVM and transformers, used to test the performance of the dataset followed by the experimental setup. We then present our classification results for in-domain model testing along with in-detailed classification error analysis and model prediction explainability. Furthermore, we evaluated the model's generalization capability across different datasets and various social media platforms.

**5.1 Classification models**

In order to design the automated classification system, we fine-tuned the state-of-the-art monolingual (Arabic) pretrained transformer architectures—AraBERT (Antoun *et al.* 2020) and QARiB (Abdelali *et al.* 2021a). Furthermore, we compared the performance of the monolingual transformers with two multilingual BERT: mBERT (Devlin *et al.* 2019) and XLM-RoBERTa (Conneau *et al.* 2019). As a baseline, we use an SVM (Platt 1998) classifier due to its reputation as a universal learner and for its ability to learn independently of the dimensionality of the feature space and the class distribution of the training dataset.

**Table 5.** Distribution of offensive and hate speech data

	Train	Dev	Test	Total
<i>OFF</i>	3172	404	887	4463
<i>NOT_OFF</i>	5716	865	1654	8235
<i>HS</i>	959	109	271	1339
<i>NOT_HS</i>	7929	1160	2270	11,359
<b>Total</b>	8888	1269	2541	12,698

### 5.1.1 Support vector machines (SVM)

We used word  $n$ -gram and character  $n$ -gram vectors weighted by term frequency-inverse document frequency (tf-idf) as features to SVM. We experimented with character  $n$ -grams ( $n = [2, 5]$ ) and word  $n$ -gram ( $n = [1, 3]$ ) individually and also a combination of them using the linear kernel with default parameters of SVM classifier ( $C = 1.0$ ).

### 5.1.2 Transformer models

We fine-tuned the mono and multilingual transformer models on our training data. For the monolingual task, both the transformers—AraBERT and QARiB—are of identical architecture with variation in the pretraining data. AraBERT is trained on Arabic Wikipedia, whereas QARiB is trained on Twitter data and Arabic Gigaword, thus encapsulating formal and informal writing styles in each transformers.

As for the multilingual models, we fine-tuned mBERT, trained on Wikipedia articles for 104 languages including Arabic using the case-sensitive base model. In addition, we also fine-tuned XLM-RoBERTa—trained on a large dataset with 100 languages using cleaned CommonCrawl data.

## 5.2 Experimental data and setup

### 5.2.1 Data split

We trained the classifiers using 70% of the data and validated with 10%. We used the rest (20% of the data) for testing the system performance. Detailed data distribution of offensive and hate speech tweets are given in Table 5.

### 5.2.2 Evaluation measures

We evaluated the dataset using macro-averaged precision, recall, and F1-measure, in addition to the accuracy.

### 5.2.3 Offensive/hate speech classifiers

For the respective downstream tasks, we fine-tuned the aforementioned BERT models by adding a dense layer that outputs class probabilities. We use a learning rate of  $8e - 5$  with a batch size of 64 and 3 epochs. For the fine-tuning, we restricted the maximum input length to 47 tokens (99th percentile in the training dataset), with no extra preprocessing of the data.

**Table 6.** Macro-averaged (P)recision, (R)ecall, and F1 for offensive language classification. Best results are highlighted in bold.

Classifier	Acc%	P	R	F1
SVM (C)	78.32	76.75	74.06	74.99
SVM (W)	72.14	69.82	70.81	70.15
SVM (C + W)	78.16	76.18	74.75	75.33
mBERT	76.43	74.09	73.32	73.66
XLM-RoBERTa*	75.00	72.50	72.47	72.48
AraBERT	82.09	80.50	79.63	80.02
QARiB	<b>84.02</b>	<b>82.53</b>	<b>82.11</b>	<b>82.31</b>

\*Results obtained with different learning rate (2e-5) instead of 8e-5. C: character *n*-grams W: words *n*-grams.

**Table 7.** Macro-averaged (P)recision, (R)ecall, and F1 for hate speech classification. Best results are highlighted in bold.

Classifier	Acc%	P	R	F1
SVM (C)	92.52	<b>84.71</b>	71.12	75.76
SVM (W)	90.08	73.90	70.73	72.15
SVM (C + W)	92.37	82.94	72.17	76.16
mBERT	91.26	77.55	73.34	75.20
XLM-RoBERTa*	92.29	79.96	78.79	79.36
AraBERT	<b>92.64</b>	81.04	<b>79.31</b>	<b>80.14</b>
QARiB	92.99	82.99	77.72	80.04

\*Results obtained with different learning rate (2e-5) instead of 8e-5. C: character *n*-grams W: words *n*-grams.

### 5.3 Results

From the reported results, presented in Table 6 and Table 7, we observe the monolingual models significantly outperform the multilingual models. This observation is in-aligned with previous studies (Polignano *et al.* 2019; Chowdhury *et al.* 2020a).

Comparing the monolingual models, for offensive classification, we observe that mixed trained model—QARiB, outperforms AraBERT by a margin of 2.3%. Contrary to that, in hate classification, we noticed AraBERT—trained with formal text—performs the best. However, the performance difference is very small (0.1%) and can be attributed to the randomness mentioned in Devlin *et al.* (2019).

#### 5.3.1 Error analysis

We analyzed the classification errors of our best classifier for offensive language detection. Common cases are listed in Table 8. We found similar errors in hate speech detection, and the main difference is that the target of the offense is a group of people as opposed to individuals.

From Table 8, we noticed most of the confusion occurs due to lack of context in the input, implicit offensive instances, misunderstanding of grammatical constructs like negation, bias

**Table 8.** Common error types in offense detection. Classes: FP (false-positives) and FN (false-negatives)

CL	Error type	Example
FP	Annotation error	الله يفشلهم. ودك تتفل بوجهه May God fail them. You wish to spit on his face
	Nonhuman target	هاشتاغ غبي. وع مره اللون مقرف Stupid hashtag. Yucky! This color is very disgusting
	Unclear context	احبيك عالفكرة بس اذا كان وجهك I salute you for the idea, but if your face
	Non-targeted offense (e.g., proverbs/idioms)	لا بارك الله للعدو. اللهم لا تؤاخذنا بما فعل السفهاء منا Oh God, do not bless the enemy (prayer)
	Animal not offensive	اتوقع رسم قصة العنز وعباله. الكلاب كلوا القطوة Expect to draw the story of the goat.. Dogs ate the cat
	Sarcasm	انا لكم الخت والام اي حركة ا اذبحكم ههه I'm your sister and mother. I will kill you haha
	FN	Culture/background
OOV/unseen words		انكتمني يابت منك ليها. خرا فيج زين Shut up you and she. Shit on you!
Need understanding		الكذب ليس سياسة تحريرية ولكن أسلوب حياة Lying is not editorial policy but a way of life
Implicit insult		حضرة الدكتور فلاح. شوف مين الي بتتكلم The doctor is a villager (uncivilized). Look who talks!
Informal writings		ما كف بعطيكك بوكسس. إنت مرت زقة I will give you a <b>blowwww</b> . You're a <b>merce nary</b>
Offensive emoji only		مش عاوز كلام كثير بس طرش له ده (الإصبع الوسطى) وبيسكت Just send (middle finger) to him and he will be silent
Negation		معد فيه حياء في البننت. قلة أدب The girl no longer has shyness. Incivility

toward some animals, and the presence of sarcasm. Moreover, we also noticed some classification errors due to human annotation errors—which is common given the complexity of the annotation task itself and its dependence on individual perspectives.

### 5.3.2 Model prediction explainability

We use LIME (Ribeiro *et al.* 2016) to interpret our best classifier. LIME is a model-agnostic method that perturbs the input text and gives weight to words based on how the model's predictions change.

We analyzed a sample of false-positive and false-negative examples from the output of LIME and our best classifier. We found much confusion in the model decisions while considering the words as offensive or clean. We summarize common errors in the examples shown in Figure 12

موتي يعطيك باللي في بالي خدى <lf> https://t.co  
 https://t.co Take <lf> . Die of your anger, you who are on  
 my mind  
 #بائع\_الكليجا خلاص الله يز عجمك خلاص rt @azo  
 rt @azo Enough! May God punish you. Enooough!  
 #Candy\_seller

**Figure 12.** Examples of explainability errors. Green-colored words contribute to classification as offensive, while red-colored words contribute to classification as a non-offensive class. Rough translations are provided.

and recommend the following for better explainability results: (i) ignore Twitter-specific symbols and symbols used for formatting (e.g., RT (Retweet), user mentions, links, <LF> (newline)), for example, replace mentions with @USER, links with URL, <LF> with space. (ii) As emojis are not recognized by LIME, replace them with their meanings, for example, middle\_finger. (iii) Remove repeated letters that are commonly found in tweets and used for emphasis, for example, convert خلاص (“xlllAS” — “enooough”) to خلاص (“xllAS” — “enough”). (iv) We noticed confusion in interpreting words’ contribution toward classifying (non-)offensiveness. However, such confusion is understandable given the limited resource that is used to learn discriminating linguistic information. In the future, we plan to explore the use of more balanced large data, collected from diverse sources. It’s worth mentioning that those errors are not necessarily LIME errors as we have errors due to model predictions (F1 is 82.31).

#### 5.4 Generalization capability of models

Our proposed method of collecting an offensive language dataset relies on emojis. While this method yields a higher percentage of offensive content compared to existing methods, it’s crucial that models trained on our dataset perform well not only on our test set but also on external datasets. Achieving such generalization can help us to assert with confidence that our method captures characteristics of offensive language that are universal, and not specific to our dataset only.

In addition, assessing the generalization of the model to other social media platforms such as Facebook and YouTube can indicate the universal usage of emoji and also the effectiveness of harnessing the extralinguistic information present in emojis for offensive tasks.

##### 5.4.1 Generalization to other Twitter data

To ascertain the generalization capability to other Twitter data, we evaluate our best performing model (QARiB, fine-tuned on our offensive training data) on the official test set of SemEval 2020 task 12 (OffensEval 2020, Arabic data) Zampieri *et al.* (2020) data. For comparison, we also fine-tune QARiB on SemEval training data and evaluate on our dataset. Further, we fine-tune QARiB on the combination of the two datasets and evaluate the test sets separately.

The SemEval data contains 8000 tweets for training and 2000 tweets for testing. The tweets are manually annotated for offensiveness. To have a higher percentage of offensive tweets, only tweets that contained “yA” in the text were considered. “yA” is commonly used in Arabic offensive language (Zampieri *et al.* 2020). Twenty percent of the data are tagged as offensive.

From Table 9, we can see that QARiB, fine-tuned on our data, achieves F1 score of 85.21 on the SemEval test set. However, when fine-tuned on SemEval data but tested on our data, QARiB achieves F1 score of 72.26. This suggests that models trained on our dataset capture characteristics

**Table 9.** Performance comparison of QARiB, fine-tuned on our dataset (EMOJI-OFF), SemEval dataset and combination of the two datasets

Train data	Test data	Acc%	P	R	F1
EMOJI-OFF	EMOJI-OFF	84.02	82.53	82.11	82.31
EMOJI-OFF	SemEval	89.30	82.40	90.05	85.21
SemEval	EMOJI-OFF	78.55	82.48	70.61	72.26
SemEval	SemEval	94.65	92.61	90.42	91.45
EMOJI-OFF + SemEval	EMOJI-OFF	83.55	81.88	81.98	81.93
EMOJI-OFF + SemEval	SemEval	94.55	91.31	91.84	91.57

**Table 10.** Performance on multi-platform data. MPOLD-TW/YT/FB refers to Twitter, YouTube, and Facebook portions of MPOLD dataset, respectively. The MPOLD-TW/YT/FB columns contain numbers reported in Chowdhury *et al.* (2020b). The EMOJI-OFF column represents QARiB fine-tuned on our data. All numbers are macro-averaged F1.

Test data	EMOJI-OFF	MPOLD-TW	MPOLD-YT	MPOLD-FB
MPOLD-TW	72.62	–	54	51
MPOLD-YT	68.2	60	–	53
MPOLD-FB	66.95	62	62	–

of SemEval data reasonably well while also having more variety compared to SemEval data. It's worth mentioning that combining the two datasets yield F1 score of 91.57 on SemEval dataset, which is an increase of 1.40 from the highest ranked system (Alami *et al.* 2020) at SemEval that achieved F1 score of 90.17 on the same test set.

#### 5.4.2 Generalization to multi-platform data

Having established generalization capability of our model to external Twitter data, we are interested in assessing how much our model can generalize to other social media platforms such as Facebook and YouTube. This is a challenging task as content on these platforms can differ greatly from each other. For example, Twitter imposes a tweet length limit of 280 characters. There is no such limitation on Facebook or YouTube comments. The audience is also different on different platforms. Because of these differences, users write in different styles across the different platforms. Moreover, as Chowdhury *et al.* (2020b) point out, the percentage of offensive content is much higher on Twitter compared to YouTube or Facebook.

To assess our model performance on other platform content, we evaluate the MPOLD dataset (Chowdhury *et al.* 2020b) using our best model (EMOJI-OFF). The MPOLD dataset is a multi-platform dataset that contains 4,000 Arabic news comments from Twitter (1624), YouTube (1592), and Facebook (784). These comments are manually tagged for offensiveness. Out of the 4,000 comments, 675 (16.88%) were tagged as offensive.

We report our cross-domain model performance in Table 10. For better comparison, we present the cross-domain results reported by Chowdhury *et al.* (2020b). From Table 10, we can see that QARiB, fine-tuned on our data, achieves F1 scores of 68.2 and 66.95 on the YouTube and Facebook portions of MPOLD dataset, respectively. While these numbers are lower than

that we saw earlier in Table 9 (as expected), they are still considerably higher than the cross-domain numbers reported in Chowdhury *et al.* (2020b), which are 0.60 and 0.62 for YouTube and Facebook test sets, respectively. This suggests that models trained on our data have the potential of generalizing to content on other platforms.

#### 5.4.3 Generalization to datasets without emoji

In order to assess how well our model performs on datasets without emoji, first, we removed all emojis from our dataset and fine-tuned our best model, QARiB. We evaluated this model on our test data and SemEval data (both with emojis removed). We did not notice any significant difference in performance as the model obtained F1 score of 85.03 on SemEval test set and F1 score of 83.05 on our test set, both of which are comparable to the numbers reported in Table 9. In the second experiment, we retained the emojis in our training set but excluded all tweets from the SemEval test set that contain emojis. Our model achieved F1 score of 83.94 in this setting. This again suggests that models trained on our data are not reliant on emojis and can generalize to datasets without emojis.

## 6. Ethics and social impact

After showcasing the effectiveness of using emojis for data collection and design, here we discuss the social impacts and ethical concerns surrounding the newly released data.

### 6.1 User privacy

The dataset was collected using the Twitter API. For privacy, we anonymized the dataset by removing usernames and other sensitive user identifiers. However, in the future, if any concern is raised for a particular content, we will comply with legitimate concerns by removing the affected tweets from the corpus. Furthermore, to keep privacy and datasets integrity, we are providing tweets by text, not the corresponding tweet ids.

### 6.2 Biases

Any biases found in the dataset are unintentional, and we do not intend to do harm to any group or individual. The bias in our data, for example, toward a particular group, is unintentional and is a true representation of the Twitter space during the period of our collection and according to our generic proposed method. Emojis used as seeds to create the collection are extracted automatically from publicly available datasets without any bias in our selection. The collection was obtained over a span of 18 months to have a good diversity of tweets written by many authors.

As for the assigned annotation labels, we follow a well-defined schema and available information to perceive final labels for offensiveness, hate speech, and vulgarity. The labels are not a reflection of our opinion. Test questions used for quality control were selected to cover different annotation classes proportionally to their distributions in the collected tweets.

Although more than 190 workers have participated in the annotation process (with at least 80% success ratio for 200 test questions), this doesn't guarantee perfect quality of annotation. As noted in Mubarak and Darwish (2014), Arabic Twitter is dominated by tweets from Gulf countries (Saudi Arabia, Kuwait, etc.) while Appen (previously CrowdFlower/Figure8) crowdsourcing platform has  $\approx 30\%$  of annotators from Egypt (the most Arabic populated country) (Mubarak and Darwish 2016). This may lead to misunderstanding (and incorrect labeling) of the portion of tweets, especially when they are written in dialects that are not understood by annotators. Moreover, almost three-quarters of Arab annotators on Appen are males which may lead to hidden biases in data annotation.

### 6.3 Potential misuse

We urge the research community to be aware that our dataset can be used to misuse like any other social media data. If such misuse is noticed, human moderation is encouraged in order to ensure this does not occur.

## 7. Conclusion

The automation of detecting offensiveness (including hate speech) is limited by the availability of large balanced manually annotated datasets. To overcome such lacking, we propose a generic language-, topic- and genre-independent approach to collect a large percentage of Arabic offensive and hate speech data. We utilize the extralinguistic information embedded in the emojis to collect a large number of offensive tweets. We manually annotated the Arabic dataset for different layers of information, including *offensiveness*, followed by fine-grained hate speech, vulgar, and violence content annotation. Furthermore, we benchmark the dataset for detecting offensive content and hate speech using different transformer architectures and performed in-depth analysis for temporal changes, linguistic usage, and target audience. We also studied specific patterns in violence tweets. We demonstrate that the transformer model trained on our data achieves strong results on an external Twitter dataset. Further, the model outperforms previous models in literature when tested on multi-platform data that contains offensive comments from Twitter, YouTube, and Facebook. This is the largest Arabic dataset for multilevel offensive language detection annotated with fine-grained hate speech, vulgar, and violence labels. We publicly release the dataset for the research community.

**Competing interests.** The authors declare none.

## References

- Abdelali A., Hassan S., Mubarak H., Darwish K. and Samih Y. (2021a). Pre-training BERT on arabic tweets: Practical considerations. *CoRR*, abs/2102.10684.
- Abdelali A., Mubarak H., Samih Y., Hassan S. and Darwish K. (2021b). QADI: Arabic dialect identification in the wild. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*. Kyiv, Ukraine (Virtual): Association for Computational Linguistics, pp. 1–10.
- Alami H., Ouatik El Alaoui S., Benlahbib A. and En-nahnahi N. (2020). LISAC FSDM-USMBA team at SemEval-2020 task 12: Overcoming AraBERT's pretrain-finetune discrepancy for Arabic offensive language identification. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*. Barcelona (online): International Committee for Computational Linguistics, pp. 2080–2085.
- Albadi N., Kurdi M. and Mishra S. (2018). Are they our brothers? Analysis and detection of religious hate speech in the arabic twittersphere. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, pp. 69–76.
- Alhazbi S. (2020). Behavior-based machine learning approaches to identify state-sponsored trolls on twitter. *IEEE Access* 8, 195132–195141.
- Alshaalan R. and Al-Khalifa H. (2020). Hate speech detection in saudi twittersphere: A deep learning approach. In *Proceedings of the Fifth Arabic Natural Language Processing Workshop*. Barcelona, Spain (Online): Association for Computational Linguistics, pp. 12–23.
- Alshehri A., El Moatez Billah Nagoudi H. A. and Abdul-Mageed M. (2018). Think before your click: Data and models for adult content in arabic twitter. In *TA-COS 2018: 2nd Workshop on Text Analytics for Cybersecurity and Online Safety*, p. 15.
- Antoun W., Baly F. and Hajj H. (2020). Arabert: Transformer-based model for arabic language understanding. In *Proceedings of The 4th Workshop on Open-Source Arabic Corpora and Processing Tools*, pp. 9–15.
- Belcastro L., Cantini R., Marozzo F., Talia D. and Trunfio P. (2020). Learning political polarization on social media using neural networks. *IEEE Access* 8, 47177–47187.
- Cheng H., Xing X., Liu X. and Lv Q. (2014). Isc: An iterative social based classifier for adult account detection on twitter. *IEEE Transactions on Knowledge and Data Engineering* 27(4), 1045–1056.

- Chowdhury S. A., Abdelali A., Darwish K., Soon-Gyo J., Salminen J. and Jansen B. J. (2020a). Improving arabic text categorization using transformer training diversification. In *Proceedings of the Fifth Arabic Natural Language Processing Workshop*, pp. 226–236.
- Chowdhury S. A., Mubarak H., Abdelali A., Jung S.-g., Jansen B. J. and Salminen J. (2020b). A multi-platform arabic news comment dataset for offensive language detection. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pp. 6203–6212.
- Chung Y.-L., Kuzmenko E., Tekiroglu S. S. and Guerini M. (2019). Conan-counter narratives through nichesourcing: A multilingual dataset of responses to fight online hate speech. *arXiv preprint arXiv: 1910*.
- Conneau A., Khandelwal K., Goyal N., Chaudhary V., Wenzek G., Guzmán F., Grave E., Ott M., Zettlemoyer L. and Stoyanov V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv: 1911*.
- Conover M., Ratkiewicz J., Francisco M. R., Gonçalves B., Menczer F. and Flammini A. (2011). Political polarization on twitter. *ICWSM 133*, 89–96.
- Darwish K., Alexandrov D., Nakov P. and Mejova Y. (2017). Seminar users in the arabic twitter sphere. In Ciampaglia G. L., Mashhadi A. and Yasseri T., (eds), *Social Informatics*. Cham: Springer International Publishing, pp. 91–108.
- Davidson T., Warmlesley D., Macy M. and Weber I. (2017). Automated hate speech detection and the problem of offensive language. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 11.
- Devlin J., Chang M.-W., Lee K. and Toutanova K. (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)*. Minneapolis, Minnesota: Association for Computational Linguistics, pp. 4171–4186.
- Dimitrov D., Ali B. B., Shaar S., Alam F., Silvestri F., Firooz H., Nakov P. and Da San Martino G. (2021a). Detecting propaganda techniques in memes. In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP'21)*.
- Dimitrov D., Ali B. B., Shaar S., Alam F., Silvestri F., Firooz H., Nakov P. and Da San Martino G. (2021b). Task 6 at semeval-2021: Detection of persuasion techniques in texts and images. In *Proceedings of the 15th International Workshop on Semantic Evaluation, SemEval*, vol. 21.
- Donato G. and Paggio P. (2017). Investigating redundancy in emoji use: Study on a twitter based corpus. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pp. 118–126.
- Dürscheid C. and Siever C. M. (2017). Beyond the alphabet–communcataion of emojis. In *Kurzfassung eines (auf Deutsch) zur Publikation eingereichten Manuskripts*.
- Gülaçtı F. (2010). The effect of perceived social support on subjective well-being. *Procedia - Social and Behavioral Sciences* 2(2), 3844–3849.
- Hassan S., Mubarak H., Abdelali A. and Darwish K. (2021a). 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*. Online. Association for Computational Linguistics, pp. 113–118.
- Hassan S., Samih Y., Mubarak H. and Abdelali A. (2020). ALT at SemEval-2020 task 12: Arabic and English offensive language identification in social media. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*. Barcelona (online): International Committee for Computational Linguistics, pp. 1891–1897.
- Hassan S., Shaar S. and Darwish K. (2021b). Cross-lingual emotion detection. *CoRR*, abs/2106.06017.
- Husain F. (2020). OSACT4 shared task on offensive language detection: Intensive preprocessing based approach. In *OSACT*, vol. 4.
- Intapong P., Charoenpit S., Achalakul T. and Ohkura M. (2017). Assessing symptoms of excessive sns usage based on user behavior and emotion: analysis of data obtained by sns apis. In *9th International Conference on Social Computing and Social Media, SCSM. 2017 held as part of the 19th International Conference on Human-Computer Interaction, HCI International 2017*. Cham: Springer, pp. 71–83.
- Kiela D., Firooz H., Mohan A., Goswami V., Singh A., Ringshia P. and Testuggine D. (2020). The hateful memes challenge: Detecting hate speech in multimodal memes. *arXiv preprint arXiv: 2005.04790*.
- Landis J. R. and Koch G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics* 33(1), 159–174.
- Mei Q. (2019). Decoding the new world language: Analyzing the popularity, roles, and utility of emojis. In *Companion Proceedings of The 2019 World Wide Web Conference*, pp. 417–418.
- Mubarak H., Abdelali A., Hassan S. and Darwish K. (2020a). Spam detection on arabic twitter. In *Social Informatics*. Cham: Springer International Publishing, pp. 237–251.
- Mubarak H. and Darwish K. (2014). Using twitter to collect a multi-dialectal corpus of Arabic. In *Proceedings of the EMNLP. 2014 Workshop on Arabic Natural Language Processing (ANLP)*, pp. 1–7.
- Mubarak H. and Darwish K. (2016). Demographic surveys of arab annotators on crowdflower. In *Weaving Relations of Trust in Crowd Work: Transparency and Reputation across Platforms*.
- Mubarak H., Darwish K. and Magdy W. (2017). Abusive language detection on arabic social media. In *Proceedings of the first workshop on abusive language online*, pp. 52–56.

- Mubarak H., Darwish K., Magdy W., Elsayed T. and Al-Khalifa H.** (2020b). Overview of osact4 arabic offensive language detection shared task. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*.
- Mubarak H., Hassan S. and Abdelali A.** (2021). Adult content detection on Arabic Twitter: Analysis and experiments. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*. Kyiv, Ukraine (Virtual): Association for Computational Linguistics, pp. 136–144.
- Mubarak H., Rashed A., Darwish K., Samih Y. and Abdelali A.** (2020c). Arabic offensive language on twitter: Analysis and experiments. *arXiv preprint arXiv: 2004.02192*.
- Nakov P., Nayak V., Dent K., Bhatwadekar A., Sarwar S. M., Hardalov M., Dinkov Y., Zlatkova D., Bouchard G. and Augenstein I.** (2021). Detecting abusive language on online platforms: A critical analysis. *arXiv preprint arXiv: 2103.00153*.
- Ousidhoum N., Lin Z., Zhang H., Song Y. and Yeung D.-Y.** (2019). Multilingual and multi-aspect hate speech analysis. *arXiv preprint arXiv: 1908.11049*.
- Platt J.** (1998). Sequential minimal optimization: A fast algorithm for training support vector machines.
- Polignano M., Basile P., de Gemmis M., Semeraro G. and Basile V.** (2019). Alberto: Italian bert language understanding model for nlp challenging tasks based on tweets. In *CLiC-it*.
- Ribeiro M. T., Singh S. and Guestrin C.** (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*. New York, NY, USA: Association for Computing Machinery, pp. 1135–1144.
- Salminen J., Hopf M., Chowdhury S. A., Jung S.-G., Almerexhi H. and Jansen B. J.** (2020). Developing an online hate classifier for multiple social media platforms. *Human-centric Computing and Information Sciences* **10**(1), 1–34.
- Waldron J.** (2012). *The Harm in Hate Speech*. Cambridge: Harvard University Press.
- Wiegand M. and Ruppenhofer J.** (2021). Exploiting emojis for abusive language detection. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 369–380.
- Zampieri M., Nakov P., Rosenthal S., Atanasova P., Karadzhov G., Mubarak H., Derczynski L., Pitenis Z. and Çöltekin C.** (2020). SemEval-2020 task 12: Multilingual offensive language identification in social media (OffensEval 2020). In *Proceedings of SemEval*.