

Meta AI at Arabic Hate Speech 2022: MultiTask Learning with Self-Correction for Hate Speech Classification

Badr AlKhamissi, Mona Diab

Responsible AI, Meta

Seattle, USA

{bkhmsi, mdiab}@fb.com

Abstract

In this paper, we tackle the Arabic Fine-Grained Hate Speech Detection shared task and demonstrate significant improvements over reported baselines for its three subtasks. The tasks are to predict if a tweet contains (1) *Offensive* language; and whether it is considered (2) *Hate Speech* or not and if so, then predict the (3) *Fine-Grained Hate Speech* label from one of six categories. Our final solution is an ensemble of models that employs multitask learning and a self-consistency correction method yielding 82.7% on the hate speech subtask—reflecting a 3.4% relative improvement compared to previous work.

1. Introduction

Disclaimer: *Due to the nature of this work, some examples contain offensiveness and hate speech. This does not reflect authors' values, however our aim is to help in detecting and preventing spread of such harmful content.*

The advent of online social networks have created a platform for billions of people to express their thoughts freely on the internet. This has enormous benefits for advancing culture. However, it also can be used by malicious actors to distribute misinformation and offensive content. This led to an increasing interest in the NLP community for the automatic detection of *Hate Speech* (HS) (Waseem and Hovy, 2016; Schmidt and Wiegand, 2017; Zampieri et al., 2019; MacAvaney et al., 2019; League, 2020; Vogels, 2021). Its dangers are becoming more apparent with studies showing its connection to hate crimes around the globe (Paz et al., 2020). Further, the spread of hateful content on the internet has also been linked to degenerate effects on peoples' psychological well-being (Gülaçti, 2010; Waldron, 2012).

HS is defined as any kind of abusive or offensive language (e.g. insults, threats, etc.) that expresses prejudice against a specific person or a group based on common characteristics such as race, religion or sexual orientation (Davidson et al., 2017; Mollas et al., 2020). Despite the growing body of HS research, few have focused on it in the context of the Arabic language.

Arabic is the mother tongue of more than 420M people, and is spoken in the fastest growing markets (Tinsley and Board, 2013). Arabic content is rapidly growing on the internet during the past couple of years (Abdelali et al., 2021). For instance, studies have shown that there are more than 27 million tweets per day in the Arab region (Alshehri et al., 2018).

In this work, we focus on the Arabic language by participating in the three subtasks of the *Arabic Fine-Grained Hate Speech Detection* shared task (Mubarak

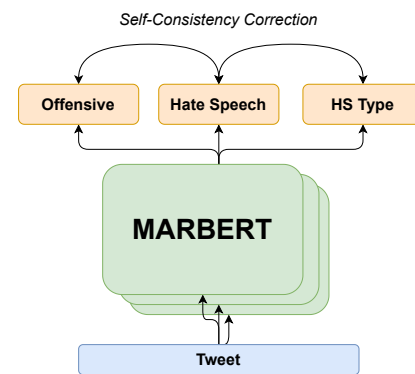


Figure 1: **System Architecture** The input tweet is encoded using a fine-tuned MARBERT model and the output embedding is given to 3 task-specific classifiers. The final prediction is computed using an ensemble of those models.

et al., 2022). The three subtasks use the same dataset from (Mubarak et al., 2022) (see Section 3 for more details.) The goal of the first subtask is to detect whether a tweet is offensive or not, while the second subtask focuses on HS detection. The third subtask further classifies a HS post into one of six fine-grained categories: Race/Ethnicity/Nationality, Religion/Belief, Ideology, Disability/Disease, Social Class, and Gender. Table 1 shows an example with its corresponding label for each subtask. An offensive post is not necessarily HS, while a HS post is always offensive. If offensive speech is not targeting an individual or a group based on common characteristics, then it is not HS.

The contributions of this paper are as follows: (1) We present a solution that outperforms the baseline models of (Mubarak et al., 2022) on every reported metric; (2) We propose the self-consistency correction method that improves the fine-grained HS subtask even further; (3) We conduct an ablation study and further analysis illustrating the importance of multi-task learning for Arabic HS detection.

Class	Example
Clean	1. لن تحصل على غدٍ افضل ما دمت تفكر بالامس (Non-Off) You won't have a better tomorrow as long you are thinking about yesterday.
Offensive	2. يلعن ابوك على هالسؤال. عساه ينقرض الكريه. (Off/Non-HS) May your father be damned for this question! I hope this fool will just wither!
Hate Speech	3. لذا القزم طلعت جايزتن له بس ماعرف يعبر. (Off/HS) This dwarf got two prizes, but he does not know how to express.

Table 1: Here we show examples and their translation in English adapted from (Mubarak et al., 2022). Example 1 is non offensive (Non-Off), Example 2 is Offensive but not Hate speech (Off/Non-HS), Example 3 is offensive and hate speech (Off/HS).

2. Motivation

There has been a growing body of research in recent years for the automatic detection of offensive language and HS online (Waseem and Hovy, 2016; Davidson et al., 2017; Schmidt and Wiegand, 2017; Fortuna and Nunes, 2018; Founta et al., 2018). Studies have shown that 41% of internet users have been harassed online with a third of these cases being targeted for something related to their inherent identity such as race or sexual orientation (Vogels, 2021; League, 2020). The massive amount of content shared on social media platforms renders manual filtering out of such malicious content impossible, driving platform providers to resort to automated means for detecting hateful content. On the other hand, machine learning based methods are data hungry and require large amounts of labelled data in order to train reliable HS classification systems. Moreover, such data has been proven hard to collect especially for low-resource languages such as Arabic. For example, (Mubarak et al., 2017) show that only 1-2% of a randomly collected sample of Arabic tweets are abusive, and only a small percentage of these are considered HS. Therefore, generalizable and robust systems for detecting offensive and HS content are direly needed.

Previous work has framed this problem as a binary classification task. However, binary judgments of HS are known to be unreliable (Sanguinetti et al., 2018; Assimakopoulos et al., 2020a). Therefore, in order to collect higher quality HS datasets researchers resorted to more complex annotation schema. For example, (Sap et al., 2020; Assimakopoulos et al., 2020b) proposed to decompose a post into several subtasks (such as the HS class and group targeted) in an effort to minimize subjectivity when deciding the HS label.

Here, we leverage the task decomposed dataset provided by (Mubarak et al., 2022) to train an Arabic transformer in a multitask manner for improving the performance of fine-grained HS detection.

3. Dataset

We use the dataset from (Mubarak et al., 2022). It consists of \sim 13k tweets in both Modern Standard Ara-

bic (MSA) and Dialectal forms of Arabic (DA). It is the largest annotated corpus of Arabic tweets that is not biased towards specific topics, genres, or dialects (Mubarak et al., 2022). Each tweet was judged by 3 annotators using crowd-sourcing. Table 2 shows the number and percentages of each annotated category. The data was split into 70% for training, 10% for development, and 20% for test. The dataset has also annotations for vulgar and violent tweets representing 1.5% and 0.7% of the whole corpus, respectively, however we are not using them in this work. Moreover, one or more user mentions are reduced to @USER, URLs are replaced with URL, and empty lines in original tweets are replaced with <LF>. See Table 1 for an example of each annotated class.

One limitation of this dataset is that the classes are highly imbalanced. Moreover, the *Disability/Disease* subclass does not exist in the training set. There are only 3 tweets related to this category and they appear in the validation and test sets only.

Class - Subclass	# of Tweets	Percentage (%)
Clean	8,235	64.85%
Offensive	4,463	35.15%
Hate Speech	1,339	10.54%
HS - Gender	641	5.05%
HS - Race	366	2.88%
HS - Ideology	190	1.50%
HS - Social Class	101	0.80%
HS - Religion	38	0.30%
HS - Disability	3	0.02%

Table 2: **Dataset Statistics** The number and percentages of tweets as represented in the entire corpus of each annotated category used in this work as described in (Mubarak et al., 2022).¹

¹The percentages do not add up to 100 since all HS tweets are a subset of the offensive ones.

Subtask	Model	Accuracy	Precision	Recall	F1 Macro
OFFD	QARiB	84.0%	82.5%	82.1%	82.3%
	AraHS	86.0%	84.6%	84.3%	84.5%
HSD	QARiB	93.0%	83.0%	77.7%	80.0%
	AraHS	94.1%	87.0%	79.5%	82.7%
HSC	AraHS	92.6%	55.1%	50.8%	51.9%

Table 3: Performance on Test for each of the subtasks Offensive Detection (OFFD), Hate Speech Detection (HSD), Hate Speech Classification (HSC) in comparison with the QARiB baseline models reported in (Mubarak et al., 2022). Our model, AraHS, outperforms the baseline QARiB on every metric. **NB:** no baseline was reported for the HSC subtask.

4. System Description

In this work, we use MARBERTv2 (Abdul-Mageed et al., 2021) as our core model. It is publicly available in the HuggingFace library (Wolf et al., 2019). It is pretrained with 1B multi-dialectal Arabic (DA) tweets which includes both MSA and DA. MARBERTv2 has the same architecture as BERT_{BASE} (Devlin et al., 2019) with ~ 163 M parameters, similarly using Word-Piece tokenization (Wu et al., 2016).

We frame the 3 subtasks as a multi-task classification problem. Specifically, the input text is encoded using MARBERTv2 and is then passed to 3 task-specific classification heads as shown in Figure 1. Each class specific head is made up of a multi-layered feed forward neural network with layer normalization (Ba et al., 2016). Concretely, the [CLS] embedding of the final MARBERTv2 transformer block is forwarded to a dense layer with 768 units, which is then passed through a GELU activation function (Hendrycks and Gimpel, 2016), the output of which is normalized using layer normalization, and this is finally given to a linear layer that maps it to the corresponding number of classes.

The final model is an ensemble of several trained models each of which uses a different set of hyperparameters. To obtain the final prediction we perform element-wise multiplication of the corresponding probabilities across the different models then take the argmax.

Self-Consistency Correction Since we are training one model for all three subtasks, and the subtasks themselves are interdependent, we leverage that to our advantage. We perform a post-processing step where errors of one classification head are corrected by the others. Concretely, the fine-grained HS prediction is corrected in the following cases: the tweet is predicted to be offensive and contains HS using the first two classification heads respectively, while the fine-grained classifier predicted that it is not HS. In that case, we take the second most probable class prediction as the label since there is an inconsistency. The other scenario in which it is corrected is when the tweet is predicted as not offensive and does not contain HS while the fine-grained classifier predicted it as one of the HS classes.

5. Experimental Setup

To train the AraHS model, we use the AdamW optimizer (Loshchilov and Hutter, 2019) and a learning rate scheduler that is warmed-up linearly for 500 steps to some initial learning rate. This is then decayed linearly to zero over the course of 10 epochs. The model is evaluated on the validation-set 4 times every epoch with equal intervals, and a checkpoint for the corresponding subtask is saved when its F1-macro score improves. The objective function is the sum of the negative log-likelihood of the three classification heads. The tokenizer encodes the input text using a maximum length of 256 tokens. The model is trained 12 times over a grid of $\{2, 4, 8, 16\}$ batch-sizes and $\{1e-5, 5e-6, 1e-6\}$ initial learning rates.

For the fine-grained HS detection subtask, we further finetune the best single model on only this subtask, using the same experimental setup described above.

6. Results

Table 3 shows the performance on the test-set for each subtask. Our method (AraHS) outperforms the baseline models reported in (Mubarak et al., 2022) on every metric: Offensive Detection subtask (OFFD) (accuracy: AraHS 86% vs. QARiB 84%; F1-Macro: AraHS 84.5% vs. QARiB 82.3%); Hate Speech Detection (HSD) subtask (accuracy: AraHS 94.1% vs. QARiB 93%; F1-Macro: AraHS 82.7% vs. QARiB 80%). Only the HS Classification (HSC) uses the self-consistency correction method. Since the dataset used in this work does not contain the disability class in the training set, the final HSC F1-macro score degrades considerably.

6.1. Ablation Study

In order to demonstrate the importance of training the subtasks jointly, we train each subtask on its own. Specifically, Table 4 compares the validation performance of each subtask with its multitask counterpart. Performance improves when using multitask learning (MTL) for the HS subtasks. However, for the offensive subtask we observe similar performance to the single-task trained models. Similar to Table 3, only the HSC is using self-consistency correction, improving the F1-macro score from 54.8% to 56.6%.

Subtask	Model	Accuracy	Precision	Recall	F1 Macro
OFFD	Single-task	88.7%	87.4%	86.1%	86.7%
	Multitask	88.5%	87.1%	86.1%	86.6%
HSD	Single-task	95.8%	87.2%	85.7%	86.4%
	Multitask	96.2%	87.7%	88.4%	88.1%
HSC	Single-task	95.3%	72.4%	46.8%	51.0%
	Multitask	95.0/94.4%	58.5/54.8%	52.5/ 58.8%	54.8/ 56.6%

Table 4: The validation performance on each subtask. The single-task models are trained on the subtask alone, while the multitask model trains all subtasks jointly. The results before and after applying the self-consistency correction as a post-processing step is shown for the HSC subtask. **Bold** shows the best result for each subtask.

7. Error Analysis

The subtasks we are training are not independent of one another, even though we are modelling them that way. For example, as previously mentioned, a tweet that is considered HS is automatically offensive as well, and needless to say that each of the fine-grained HS classes (e.g. race, gender, etc.) are HS. Therefore, in this section we want to measure the degree of self-consistency within the trained multitask model. Specifically, we take the predictions of the best ensemble of models and calculate the percentage of contradiction between each classification head (see Table 5). Concretely, we compute the number of times in which the OFFD head yielded a negative prediction whereas the HSD or the HSC yielded a positive one. This is a contradiction since each HS post must be offensive as well. Similarly, we calculate the number of times the HSD head predicted negative while the HSC predicted positive and vice versa. As illustrated in Table 5, correcting the HSC head based on the two other subtasks reduces the contradiction considerably (from 2.6% to 0.79%) while achieving a better performance overall.

Subtask	Contradiction (%)
OFFD	2.44%
HSD	2.60%
HSC	2.60% / 0.79%

Table 5: Percentage of contradiction between classification heads before and after self-correction for the HSC. Each row correspond to the best checkpoint achieved for that subtask.

Furthermore, Table 6 shows two examples where the classification heads disagreed with one another. For example, the first tweet was detected as HS but the fine-grained classification head classified it as non-HS leading to a disagreement. Using our self-consistency correction method, the model was able to correct itself and yield the correct label, which was the *Ideology* subclass in this case. Example 1 in the table is a modification of

an Arabic adage: “الطيور على اشكالها تقع”, corresponding to “Birds of a feather flock together. Changing *birds* to *frogs* implies ugliness. Such tweets are not straight forward to classify since they require an understanding of cultural knowledge and implicit social nuance that is not explicitly encoded in language models such as MARBERT. One way to mitigate this is to finetune the model on a corpus that contains such information explicitly incorporating such inductive bias. Another method would be training the language model to generate the implication of the tweet as an additional subtask. The other example in the table implies that people of a certain nationality are ignorant. We believe that the provided gold label is incorrect (not HS). We believe that this tweet constitutes HS because it is offensive (a certain group of people is ignorant since they parrot rather than understand information) and it targets a group. Accordingly, the model was able to successfully predict it as HS, and even yield the correct class for it using the self-consistency method. In Table 7 we report the percentage of false positives (FP) and false negatives (FN) of the best checkpoint of each subtask. To compute the percentage of FP and FN for the HSC subtask we convert it into a binary variable with negative implying that the prediction is not-HS and positive otherwise. Interestingly, the self-consistency correction method increases the percentage of FPs, as it takes the second top prediction as its label when both the HSD and OFFD are positive. We note that HS systems can tolerate more false positives (i.e. over enforcement) than false negatives (i.e. under enforcement), since the latter will lead to more propagation of harm. This highlights the advantage of self-consistency correction.

8. Related Work

Datasets The first Arabic HS dataset was collected by (Albadi et al., 2018) and consisted of $\sim 6.6k$ Arabic HS tweets. In an effort to collect a more dialect specific dataset, (Haddad et al., 2019) compiled 6k tweets of the Tunisian dialect containing both abusive language and HS. (Mulki et al., 2019) similarly collected a Levantine HS dataset. In the multilingual front, (Ousidhoum et al., 2019) created a HS dataset made up of 13k Arabic, English and French tweets with fine-grained labels

Tweet	OFFD	HSD	HSC
الضفادع على اشكالها تقع Frogs settle with their own kind.	✓—✓	✓—✓	✗— Ideology
أستغفر الله العظيم شعب حافظ مش فاهم The people of this country memorize without understanding.	✓—✗	✓—✗	✗—✗
ضروري من الطقوس في هذا المجتمع القائم على النفاق الاجتماعي Necessary rituals in this community that is based on social hypocrisy.	✗—✓	✗—✓	✗— Nationality

Table 6: The first two are examples that led to a disagreement between the classification heads, while the third one show a false negative example. Below each example is a rough English translation that is used to convey the meaning. On the right-side of the table, the prediction of the model is shown first followed by the ground truth label for each subtask. Note that the self-consistency correction method was able to correct the first two examples among others.

	False +ve (%)	False -ve (%)
OFFD	4.96%	6.54%
HSD	1.97%	1.81%
HSC	2.52%/3.86%	2.44%/1.73%

Table 7: Percentage of false positives and false negatives for the best checkpoint for each subtask. In the HSC we report the percentage before and after applying the self-consistency correction method.

covering different aspects such as target groups, directness, target attributes and hostility types.

Models Early work tackled this problem by extracting n-gram features using term frequency weighting, which was then passed to a Support Vector Machine (SVM) and Naive Bayes (NB) classifiers (Mulki et al., 2019). Other work used a gated recurrent unit (GRU) coupled with an SVM trained on the AraVec embeddings (Ashi et al., 2018) to classify HS (Albadi et al., 2018). (Hassan et al., 2020) used an ensemble of SVM, CNN-BiLSTM and feed-forward neural networks for HS detection. (Duwairi et al., 2021) showed that CNN models outperform their CNN-LSTM and CNN-BiLSTM counterparts in detecting HS when treated as a binary classification task on the ArHS dataset.

Multitask Learning In the Offensive Detection shared task co-located with the 4th Workshop on Open-Source Arabic Corpora and Processing Tools (OS-ACT4) (Al-Khalifa et al., 2020), (Djandji et al., 2020) trained AraBERT (Antoun et al., 2020) on multiple tasks simultaneously achieving the best score on the shared task. (El Mahdaouy et al., 2021) used a model based on MARBERT that employed MTL. In another line of work, a CNN-BiLSTM based architecture was trained using MTL to detect HS and OFF language (Abu Farha and Magdy, 2020). That model used extra sentiment information (Abu Farha and Magdy, 2019) during training. Finally, (Aldjanabi et al., 2021) explore a multi-corpus-based learning approach built on top of MARBERT. It uses MTL from three datasets

for improving OFF and HS detection. Unlike previous work, our paper focuses on improving fine-grained Arabic HS classification using MTL and self-consistency correction on the new dataset introduced in (Mubarak et al., 2022).

9. Conclusion

In this paper, we propose MTL as an approach to Hate Speech Classification. Our proposed model, AraHS, outperforms the baseline models. AraHS is an ensemble of MARBERT (Abdul-Mageed et al., 2021) models trained with different hyperparameters using MTL. The fine-grained HS subtask is then finetuned on its own for a couple of epochs. We demonstrate the importance of training the three subtasks jointly through an ablation study and propose the self-consistency correction method that improves the final result even further. In future work, we would like to explore the limits of combining multilingual models (e.g. mBART (Liu et al., 2020)) with Arabic monolingual models such as MARBERT. Further, we would like to explore treating the problem as a conditional generation task using the AraT5 model (Nagoudi et al., 2021) that has been shown to outperform MARBERT on the Arabic language understanding evaluation benchmark (ARLUE) (Abdul-Mageed et al., 2021).

10. Ethics and Social Impact

Modern deep learning models are energy intensive and can cause environmental damage due to the carbon dioxide emissions required for running modern hardware. Studies have shown that training a BERT model on GPU has a comparable carbon footprint to a trans-American flight (Strubell et al., 2019). In this work, even though we do not pre-train the model, we still run multiple experiments across a grid of hyperparameters, that when combined consumes significant energy. Therefore, one of the reasons we chose multitask learning (MTL) is that we can reduce the amount of training substantially by training one model on multiple tasks. MTL does not only offer energy efficiency, but is

also more data efficient, it has been shown to converge faster by leveraging auxiliary information and reduces over-fitting through shared representations (Crawshaw, 2020).

Further, building models for detecting OFF language and HS can help improve the moderation of hateful content on the internet. This can potentially lead to less hate crimes and better psychological well-being for users receiving such content. However, the authors are aware of potential misuse of HS models, such as propagating the spread of HS rather than suppressing it. Therefore, human moderation is required for preventing such misuse.

11. Bibliographical References

- Abdelali, A., Mubarak, H., Samih, Y., Hassan, S., and Darwish, K. (2021). QADI: Arabic dialect identification in the wild. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 1–10, Kyiv, Ukraine (Virtual), April. Association for Computational Linguistics.
- Abdul-Mageed, M., Elmadany, A., and Nagoudi, E. M. B. (2021). ARBERT & MARBERT: Deep bidirectional transformers for Arabic. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7088–7105, Online, August. Association for Computational Linguistics.
- Abu Farha, I. and Magdy, W. (2019). Mazajak: An online Arabic sentiment analyser. In *Proceedings of the Fourth Arabic Natural Language Processing Workshop*, pages 192–198, Florence, Italy, August. Association for Computational Linguistics.
- Abu Farha, I. and Magdy, W. (2020). Multitask learning for Arabic offensive language and hate-speech detection. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 86–90, Marseille, France, May. European Language Resource Association.
- Al-Khalifa, H., Magdy, W., Darwish, K., Elsayed, T., and Mubarak, H. (2020). Proceedings of the 4th workshop on open-source arabic corpora and processing tools, with a shared task on offensive language detection. Marseille, France, May. European Language Resource Association.
- Albadi, N., Kurdi, M., and Mishra, S. (2018). Are they our brothers? analysis and detection of religious hate speech in the arabic twittersphere. *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 69–76.
- Alldjanabi, W., Dahou, A., Al-qaness, M. A. A., Elsayed Abd Elaziz, M., Helmi, A., and Damasevicius, R. (2021). Arabic offensive and hate speech detection using a cross-corpora multi-task learning model. *Informatics*, 8:69, 10.
- Alshehri, A., Nagoudi, E. M. B., Alhuzali, H., and Abdul-Mageed, M. (2018). Think before your click: Data and models for adult content in arabic twitter.
- 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, with a Shared Task on Offensive Language Detection*, pages 9–15, Marseille, France, May. European Language Resource Association.
- Ashi, M. M., Siddiqui, M. A., and Nadeem, F. (2018). Pre-trained word embeddings for arabic aspect-based sentiment analysis of airline tweets. In *ASIS*.
- Assimakopoulos, S., Muskat, R. V., van der Plas, L., and Gatt, A. (2020a). Annotating for hate speech: The maneco corpus and some input from critical discourse analysis. *arXiv preprint arXiv:2008.06222*.
- Assimakopoulos, S., Vella Muskat, R., van der Plas, L., and Gatt, A. (2020b). Annotating for hate speech: The MaNeCo corpus and some input from critical discourse analysis. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5088–5097, Marseille, France, May. European Language Resources Association.
- Ba, J., Kiros, J., and Hinton, G. E. (2016). Layer normalization. *ArXiv*, abs/1607.06450.
- Crawshaw, M. (2020). Multi-task learning with deep neural networks: A survey. *ArXiv*, abs/2009.09796.
- Davidson, T., Warmsley, D., Macy, M. W., and Weber, I. (2017). Automated hate speech detection and the problem of offensive language. In *ICWSM*.
- 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)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Djandji, M., Baly, F., Antoun, W., and Hajj, H. (2020). Multi-task learning using AraBert for offensive language detection. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 97–101, Marseille, France, May. European Language Resource Association.
- Duwairi, R., Hayajneh, A., and Quwaider, M. (2021). A deep learning framework for automatic detection of hate speech embedded in arabic tweets. *Arabian Journal for Science and Engineering*, 46:1–14.
- El Mahdaouy, A., El Mekki, A., Essefar, K., El Mamoun, N., Berrada, I., and Khoumsi, A. (2021). Deep multi-task model for sarcasm detection and sentiment analysis in Arabic language. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 334–339, Kyiv,

- Ukraine (Virtual), April. Association for Computational Linguistics.
- Fortuna, P. and Nunes, S. (2018). A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30.
- Founta, A. M., Djouvas, C., Chatzakou, D., Leontiadis, I., Blackburn, J., Stringhini, G., Vakali, A., Sirivianos, M., and Kourtellis, N. (2018). Large scale crowdsourcing and characterization of twitter abusive behavior. In *Twelfth International AAAI Conference on Web and Social Media*.
- Gülaçtı, F. (2010). The effect of perceived social support on subjective well-being. *Procedia - Social and Behavioral Sciences*, 2:3844–3849.
- Haddad, H., Mulki, H., and Oueslati, A., (2019). *T-HSAB: A Tunisian Hate Speech and Abusive Dataset*, pages 251–263. 10.
- Hassan, S., Samih, Y., Mubarak, H., Abdelali, A., Rashed, A., and Chowdhury, S. A. (2020). ALT submission for OSACT shared task on offensive language detection. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 61–65, Marseille, France, May. European Language Resource Association.
- Hendrycks, D. and Gimpel, K. (2016). Bridging nonlinearities and stochastic regularizers with gaussian error linear units. *ArXiv*, abs/1606.08415.
- League, A.-D. (2020). Online hate and harassment. the american experience 2021. *Center for Technology and Society*. Retrieved from www.adl.org/media/14643/download.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., and Zettlemoyer, L. (2020). Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Loshchilov, I. and Hutter, F. (2019). Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- MacAvaney, S., Yao, H.-R., Yang, E., Russell, K., Goharian, N., and Frieder, O. (2019). Hate speech detection: Challenges and solutions. *PloS one*, 14(8):e0221152.
- Mollas, I., Chrysopoulou, Z., Karlos, S., and Tsoumakas, G. (2020). Ethos: an online hate speech detection dataset. *ArXiv*, abs/2006.08328.
- 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*, pages 52–56, Vancouver, BC, Canada, August. Association for Computational Linguistics.
- Mubarak, H., Hassan, S., and Chowdhury, S. A. (2022). Emojis as anchors to detect arabic offensive language and hate speech. *arXiv preprint arXiv:2201.06723*.
- Mulki, H., Haddad, H., Bechikh Ali, C., and Alshabani, H. (2019). L-HSAB: A Levantine Twitter dataset for hate speech and abusive language. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 111–118, Florence, Italy, August. Association for Computational Linguistics.
- Nagoudi, E. M. B., Elmadany, A., and Abdul-Mageed, M. (2021). Arat5: Text-to-text transformers for arabic language generation.
- Ousidhoum, N., Lin, Z., Zhang, H., Song, Y., and Yeung, D.-Y. (2019). Multilingual and multi-aspect hate speech analysis. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4675–4684, Hong Kong, China, November. Association for Computational Linguistics.
- Paz, M. A., Montero-Díaz, J., and Moreno-Delgado, A. (2020). Hate speech: A systematized review. *SAGE Open*, 10(4):2158244020973022.
- Sanguinetti, M., Poletto, F., Bosco, C., Patti, V., and Stranisci, M. (2018). An Italian Twitter corpus of hate speech against immigrants. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Sap, M., Gabriel, S., Qin, L., Jurafsky, D., Smith, N. A., and Choi, Y. (2020). Social bias frames: Reasoning about social and power implications of language. In *ACL*.
- Schmidt, A. and Wiegand, M. (2017). A survey on hate speech detection using natural language processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10, Valencia, Spain, April. Association for Computational Linguistics.
- Strubell, E., Ganesh, A., and McCallum, A. (2019). Energy and policy considerations for deep learning in nlp. *ArXiv*, abs/1906.02243.
- Tinsley, T. and Board, K. (2013). *Languages for the Future*. British Council.
- Vogels, E. A. (2021). The state of online harassment. *Pew Research Center*, 13.
- Waldron, J. (2012). *The Harm in Hate Speech*. Harvard University Press.
- Waseem, Z. and Hovy, D. (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, San Diego, California, June. Association for Computational Linguistics.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., and Brew, J. (2019). Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Wu, Y., Schuster, M., Chen, Z., Le, Q., Norouzi,

M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, u., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., and Dean, J. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. 09.

Zampieri, M., Malmasi, S., Nakov, P., Rosenthal, S., Farra, N., and Kumar, R. (2019). Predicting the type and target of offensive posts in social media. 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 1415–1420, Minneapolis, Minnesota, June. Association for Computational Linguistics.