

Leveraging Affective Bidirectional Transformers for Offensive Language Detection

AbdelRahim Elmadany, Chiyu Zhang, Muhammad Abdul-Mageed, Azadeh Hashemi

{a.elmadany, muhammad.mageeed, azadeh.hashemi}@ubc.ca, chiyuzh@mail.ubc.ca

Natural Language Processing Lab
University of British Columbia

Abstract

Social media are pervasive in our life, making it necessary to ensure safe online experiences by detecting and removing offensive and hate speech. In this work, we report our submission to the Offensive Language and hate-speech Detection shared task organized with the 4th Workshop on Open-Source Arabic Corpora and Processing Tools Arabic (OSACT4). We focus on developing purely deep learning systems, without a need for feature engineering. For that purpose, we develop an effective method for automatic data augmentation and show the utility of training both offensive and hate speech models off (i.e., by fine-tuning) previously trained affective models (i.e., sentiment and emotion). Our best models are significantly better than a vanilla BERT model, with 89.60% acc (82.31% macro F_1) for hate speech and 95.20% acc (70.51% macro F_1) on official TEST data.

1. Introduction

Social media are widely used at a global scale. Communication between users from different backgrounds, ideologies, preferences, political orientations, etc. on these platforms can result in tensions and use of offensive and hateful speech. This negative content can be very harmful, sometimes with real-world consequences. For these reasons, it is desirable to control this type of uncivil language behavior by detecting and removing this destructive content.

Although there have been a number of works on detecting offensive and hateful content in English (e.g. (Agrawal and Awekar, 2018; Badjatiya et al., 2017; Nobata et al., 2016)), works on many other languages are either lacking or rare. This is the case for Arabic, where there have been only very few works (e.g., (Alakrot et al., 2018; Albadi et al., 2018; Mubarak et al., 2017; Mubarak and Darwish, 2019)). For these motivations, we participated in the Offensive Language and hate-speech Detection shared task organized with the 4th Workshop on Open-Source Arabic Corpora and Processing Tools Arabic (OSACT4).

Offensive content and hate speech are less frequent online than civil, acceptable communication. For example, only 19% and $\sim 5\%$ of the released shared task data are offensive and hate speech, respectively. This is the case in spite of the fact that the data seems to have been collected based on trigger seeds that are more likely to accompany this type of harmful content. As such, it is not easy to acquire data for training machine learning systems. For this reason, we direct part of our efforts to automatically augmenting training data released by the shared task organizers (Section 3.1.). Our experiments show the utility of our data enrichment method. In addition, we hypothesize trained affective models can have useful representations that might be effective for the purpose of detecting offensive and hateful content. To test this hypothesis, we fine-tune one sentiment analysis model and one emotion detection model on our training data. Our experiments support our hypothesis (Section 4.). All our models are based on the Bidirectional

Encoder from Transformers (BERT) model. Our best models are significantly better than competitive baseline based on vanilla BERT. Our contributions can be summarized as follows:

- We present an effective method for automatically augmenting training data. Our method is simple and yields sizable additional data when we run it on a large in-house collection.
- We demonstrate the utility of fine-tuning off-the-shelf affective models on the two downstream tasks of offensive and hate speech.
- We develop highly accurate deep learning models for the two tasks of offensive content and hate speech detection.

The rest of the paper is organized as follows: We introduce related works in Section 2., shared task data and our datasets in Section 3., our models in Section 4., and we conclude in Section 5..

2. Related Work

Thematic Focus: Research on undesirable content shows that social media users sometimes utilize *profane*, *obscene*, or *offensive* language (Jay and Janschewitz, 2008; Wiegand et al., 2018); *aggression* (Kumar et al., 2018; Modha et al., 2018); *toxic content* (Georgakopoulos et al., 2018; Fortuna et al., 2018; Zampieri et al., 2019), and *bullying* (Dadvar et al., 2013; Agrawal and Awekar, 2018; Fortuna et al., 2018).

Overarching Applications: Several works have taken as their target detecting these types of negative content with a goal to build applications for (1) content filtering or (2) quantifying the intensity of polarization (Barberá and Sood, 2015; Conover et al., 2011), (3) classifying trolls and propaganda accounts that often use offensive language (Darwish et al., 2017), (4) identifying hate speech that may correlate with hate crimes (Nobata et al., 2016), and (5) detecting signals of conflict, which are often

	Dataset	#tweets	# NOT_OFF	# OFF	OFF%	# NOT_HS	# HS	HS %
Shard-task data	TRAIN	6994	5585	1409	20%	6633	361	5%
	DEV	1000	821	179	18%	956	44	4%
	TEST	2000	-	-	-	-	-	-
Augmented data	AUG-TRAIN-HS	209780	-	-	-	199291	10489	5%
	AUG-TRAIN-OFF	480777	215365	265413	55%	-	-	-

Table 1: Offensive (OFF) and Hate Speech (HS) Labels distribution in datasets

preceded by verbal hostility (Chadefaux, 2014).

Methods: A manual way for detecting negative language can involve building a list of offensive words and then filtering text based on these words. As Mubarak and Darwish (2019) also point out, this approach is limited because (1) offensive words are ever evolving with new words continuously emerging, complicating the maintenance of such lists and (2) the offensiveness of certain words is highly context- and genre-dependent and hence a lexicon-based approach will not be very precise. Machine learning approaches, as such, are much more desirable since they are more nuanced to domain and also usually render more accurate, context-sensitive predictions. This is especially the case if there are enough data to train these systems.

Most work based on machine learning employs a supervised approach at either (1) character level (Malmasi and Zampieri, 2017), (2) word level (Kwok and Wang, 2013), or (3) simply employ some representation incorporating word embeddings (Malmasi and Zampieri, 2017). These studies use different learning methods, including Naive Bayes (Kwok and Wang, 2013), SVMs (Malmasi and Zampieri, 2017), and classical deep learning such as CNNs and RNNs (Nobata et al., 2016; Badjatiya et al., 2017; Alakrot et al., 2018; Agrawal and Awekar, 2018). Accuracy of the aforementioned systems range between 76% and 90%. It is also worth noting that some earlier works (Weber et al., 2013) use sentiment words as features to augment other contextual features. Our work has affinity to this last category since we also leverage affective models trained on sentiment or emotion tasks. Our approach, however, differs in that we build models free of hand-crafted features. In other words, we let the model learn its representation based on training data. This is a characteristic attribute of deep learning models in general.¹ In terms of the specific information encoded in classifiers, researchers use profile information in addition to text-based features. For example, Abozinadah (2017) apply SVMs on 31 features extracted from user profiles in addition to social graph centrality measures.

Methodologically, our work differs in three ways: (1) we train offensive and hate speech models off affective models (i.e., we fine-tune already trained sentiment and emotion models on both the offensive and hate speech tasks). (2)

¹Of course hand-crafted features can also be added to a representation fed into a deep learning model. However, we do not do this here.

We apply BERT language models on these two tasks. We also (3) automatically augment offensive and hate speech training data using a simple data enrichment method.

Arabic Offensive Content: Very few works have been applied to the Arabic language, focusing on detecting offensive language. For example, (Mubarak et al., 2017) develop a list of obscene words and hashtags using patterns common in offensive and rude communications to label a dataset of 1,100 tweets. Mubarak and Darwish (2019) applied character n-gram FasText model on a large dataset (3.3M tweets) of offensive content. Our work is similar to Mubarak and Darwish (2019) in that we also automatically augment training data based on an initial seed lexicon.

3. Data

In our experiments, we use two types of data: (1) data distributed by the Offensive Language Detection shared task and (2) an automatically collected dataset that we develop (Section 3.1.). The shared task dataset comprises 10,000 tweets manually annotated for two sub-tasks: *offensiveness* (Sub_task_A)² and *hate speech* (Sub_task_B)³. According to shared task organizers,⁴ offensive tweets in the data contain explicit or implicit insults or attacks against other people, or inappropriate language. Organizers also maintain that hate speech tweets contains insults or threats targeting a specific group of people based on the nationality, ethnicity, gender, political or sport affiliation, religious belief, or other common characteristics of such a group. The dataset is split by shared task organizers into 70% TRAIN, 10% DEV, and 20% TEST. Both labeled TRAIN and DEV splits were shared with participating teams, while tweets of TEST data (without labels) was only released briefly before competition deadline.

It is noteworthy that the dataset is imbalanced. For *offensiveness* (Sub_task_A), only 20% of the TRAIN split are labeled as offensive and the rest is not offensive. For *hate speech* (Sub_task_B), only 5% of the tweets are annotated as hateful. Due to this imbalanced, the official evaluation metric is macro F₁ score. Table 1 shows the size and label distribution in the shared task data splits.⁵

²<https://competitions.codalab.org/competitions/22825>.

³<https://competitions.codalab.org/competitions/22826>

⁴<http://edinburghnlp.inf.ed.ac.uk/workshops/OSACT4/>.

⁵Table 1 also shows size and class distribution for our auto-

The following are example tweets from the shared task TRAIN split.

Examples of offensive and hateful tweets:

- 1) يا رب يا واحد يا أحد بحق يوم الاحد ان
تهلك بني سعود المجرمين. لاجل اطفال اليمن شاركوا.
*Oh my Lord, O One and Only, destroy the family of
Sau'd, for they are the criminals who put children of
Yemen to suffer.*⁶
- 2) يا لبناني يا فضلات الاستعمار الفرنسي اللبنانيين
بالخليج يشغلون نسوانهم عاهرات
*Hey, you Lebanese guy, you're the wastes of the
French colonizers. The Lebanese in the Gulf put their
women in prostitution work.*

Examples for offensive but not hate-speech tweets:

- 3) يا لطيف.. يا ساتر .أحمدوا ربكم إنها
مستعدة كيف لو إنها واقفه
*Oh my lord... Thank God she has disability. What
would have happened if she were not disabled?*
- 4) يا ترى مخيلنا ايه يا چون سنو انت
و الكلوبية اللي جنبك دي
*I wonder what you, and this little pitch by your side,
are hiding for us, John Snow?*

Examples for not offensive and not hate-speech tweets:

- 5) يا بكون بحياتك الأهم يا إما ما بدني أكون
*Either I become the most important in your life, or I
become nothing at all.*
- 6) ايش الاكل الحلو ذا يا شميه تسلم يدك ياعسل
ياقشطه ياحلوه ياشكره يا طباخه يا فنانه يا كل شي
*Wow! How wonderful this food is, Sumaia! You're
such a honey, beauty, sweetie, and good cook! You're
are artist! You're everything!*

3.1. Data Augmentation

As explained earlier, the positive class in the offensive sub-task (i.e., the category 'offensive') is only 20% and in the hateful sub-task (i.e., the class 'hateful') it is only 5%. Since our goal is to develop exclusively deep learning

automatically extracted dataset, to which we refer to as augmented (AUG).

⁶Original tweets can be run-on sentences, lack proper grammatical structures or punctuation. In presented translation, for readability, while we maintain the meaning as much as possible, we render grammatical, well-structured sentence.

models, we needed to extend our training data such that we increase the positive samples. For this reason, we develop a simple method to automatically augment our training data. Our method first depends on extracting tweets that contain any of a seed lexicon (explained below) and satisfy a predicted sentiment label condition. We hypothesize that both offensive and hateful content would carry negative sentiment and so it would be intuitive to restrict any automatically extracted tweets to those that carry these negative sentiment labels. To further test this hypothesis, we analyzing the distribution of the sentiment classes in the TRAIN split using an off-the-shelf tool, AraNet (Abdul-Mageed et al., 2020). As shown in Figure 3, AraNet assigns sensible sentiment labels to the data. For the 'offensive' class, the tool assigns 65% negative sentiment tags and for the non-offensive class it assigns only 60% positive sentiment labels.⁷ For the hate speech data, we find that AraNet assigns 72% negative labels to the 'hateful' class and 55% positive sentiment labels for the 'non-hateful' class. Based on this analysis, we decide to impose a sentiment-label condition on the automatically extended data as explained earlier. In other words, we only choose 'offensive' and 'hateful' class data from tweets predicted as negative sentiment. Similarly, we only choose 'non-offensive' and 'non-hateful' tweets assigned positive sentiment labels by AraNet. We now explain how we extend the dataset. We now explain our approach to extract tweets with an offensive and hateful seed lexicon.

To generate a seed lexicon, we extract all words that follow the *Ya* (*Oh, you*) in the shared task TRAIN split positive class in the two sub-tasks (i.e., 'offensive' and 'hateful'). The intuition here is that the word *Ya* acts as a trigger word that is likely to be followed by negative lexica. This gives us a set of 2,158. We find that this set can have words that are neither offensive nor hateful outside context and so we manually select a smaller set of 352 words that we believe are much more likely to be effective offensive seeds and only 38 words that we judge as more suitable carriers of hateful content. Table 2 shows samples of the offensive and hateful seeds. Table 3 shows examples of seeds in our initial larger set that we filtered out since these are less likely to carry negative meaning (whether offensive or hateful).

To extend the offensive and hateful tweets, we use 500K randomly sampled, unlabeled, tweets from (Abdul-Mageed et al., 2019) that each have at least one occurrence of the trigger word *Ya* and at least one occurrence of a word from either of our two seed lexica (i.e., the offensive and hateful seeds).⁸ We then apply AraNet (Abdul-Mageed et al., 2020) on this 500K collection and keep only tweets assigned negative sentiment labels. Tweets that

⁷AraNet (Abdul-Mageed et al., 2020) assigns only positive and negative sentiment labels. In other words, it does not assign neutral labels.

⁸The 500K collection is extracted via searching a larger sample of $\sim 21M$ tweets that all have the trigger word *Ya*. This corpus is also taken from (Abdul-Mageed et al., 2019). Note that a tweet can have both an offensive and a hateful seed.

Arabic Offensive	English	Arabic Hateful	English
يا خبيث	You, fat ass!	يا منجاوي	You're Manjawi
يا رعا	You're mobby	يا دندراوي	You're Dandarawi
يا متشرد	You're a tramp	يا سعوديين	You Saudis
يا مجانين	You're crazy	يا دحباشي	You're Dahbashi
يا جعان	You, hungry man!	يا ادعائي	You, false claimer
يا فاجره	You, morally loose	يا حوثي	You, Houthi
يا شمال	Oh, whore	يا شيعي	You, Shiite
يا زباله	You, junky	يا عميل	You, spay
يا بهائم	You, animals	يا اخوانحي	You, Ikhwangis
يا بغض	You, hateful	يا اخوان	You, Ikhwan
يا وسخه	You, dirty woman	يا ابن الحرابيع	You, son of tramps
يا طاغيه	You, tyrant	يا ابن الحرام	You, bastard
يا فاجر	You, salacious	يا ابن الزنا	You, bastard
يا مغفل	You, idiot	يا ابن اليهوديه	You, son of Jewish woman
يا رخمه	You, silly woman	يا ابن الزانيه	You, son of adulterous woman
يا نحس	You, sinister	يا ابن الموهومه	You, son of deceived woman
يا غباء	You, stupid head	يا ابن العرص	You, son of pimp
يا كئيب	You, gloomy head	يا ابن المتناكه	You, son of adulterous
يا مرا	You, unworthy woman	يا امارات	You, Emirate
يا حقمي	You, fools	يا اتحادي	You, Itihadi

Table 2: Examples of offensive and hateful seeds in our lexica

carry offensive seeds are labeled as ‘offensive’ and those carrying hateful seeds are tagged as ‘hateful’. This gives us 265,413 offensive tweets and 10,489 hateful tweets. For reference, the majority (%=67) of the collection extracted with our seed lexicon are assigned negative sentiment labels by AraNet. This reflects the effectiveness of our lexicon as it matches our observations about the distribution of sentiment labels in the shared task TRAIN split.

To add positive class data (i.e., ‘not-offensive’ and ‘not-hateful’) to this augmented collection, we randomly sample another 500K tweets that carry *Ya* from (Abdul-Mageed et al., 2019) that do not carry any of the two offensive and hateful seed lexica. We apply AraNet on these tweets and keep only tweets assigned a positive sentiment label (%=70). We use 215,365 tweets as ‘non-offensive’ but only 199,291 as ‘non-hateful’.⁹ Table 1 shows the size and distribution of class labels in our extended dataset.

Figure 2 and Figure 1 are word clouds of unigrams in our extended training data (offensive and hateful speech, respectively) after we remove our seed lexica from the data. The clouds show that the data carries lexical cues likely to occur in each of the two classes (offensive and hateful). Examples of frequent words in the offensive class include *dog*, *animal*, *son of*, *mother*, *dirty woman*, *monster*, *mad*, and *on you*. Examples in the hateful data include *shut up*, *dogs*, *son of*, *animal*, *dog*, *haha*, and *for this reason*. We note that the

hateful words do not include direct names of groups since these were primarily our seeds that we removed before we prepare the word cloud. Overall, the clouds provide sensible cues of our phenomena of interest across the two tasks.



Figure 1: A word cloud of unigrams in our extended training offensive data (AUG-TRAIN-OFF).

4. Models

4.1. Data Pre-Processing

We perform light Twitter-specific data cleaning (e.g., replacing numbers, usernames, hashtags, and hyperlinks by unique tokens NUM, USER, HASH, and URL respectively). We also perform Arabic-specific normalization (e.g., removing diacritics and mapping various forms of Alef and Yeh each to a canonical form). For text tokenization, we use byte-pair encoding (PBE) as implemented in

⁹We decided to keep only 199,291 ‘non-hateful’ tweets since our augmented ‘hateful’ class comprises only 10,489 tweets.

Arabic Non-OFF/Non-HS	English
يا ابطال	You, heros
يا فنان	You, artist
يا عبسميع	You, Absemee'
يا عالم	Oh, people
يا منسي	Oh, forgotten man
يا ما عندك	You have a lot
يا ماما	Oh, mum
يا قمرر	You, beautiful lady
يا جامد	You, wonderful man
يا طفل	Oh, child
يا قطه	Oh, delicate lady
يا حيلتها	You, lulled
يا ايتها	Oh you. . .
يا راعي	You, caregiver
يا حبيبي	Oh, darling
يا رب	Oh, my Lord
يا واحد	Oh, the One
يا ناس	You, people
يا اخر	Oh, the Last
يا بابا	Oh, daddy

Table 3: Examples of non-offensive/non-hateful seeds filtered out from our lexica.



Figure 2: A word cloud of unigrams in our extended training hate speech data (AUG-TRAIN-HS).

Multilingual Cased BERT model.

4.2. BERT

Our experiments are based on BERT-Base Multilingual Cased model released by (Devlin et al., 2018)¹⁰. BERT stands for **B**idirectional **E**ncoder **R**epresentations from **T**ransformers. It is an approach for language modeling that involves two self-supervised learning tasks, (1) masked language models (MLM) and (2) next sentence predication (NSP). BERT is equipped with an Encoder architecture which naturally conditions on bi-directional context. It randomly masks a given percentage of input tokens and attempts to predict these masked tokens. (Devlin et al., 2018) mask 15% of the tokens (the authors use

¹⁰<https://github.com/google-research/bert/blob/master/multilingual.md>.

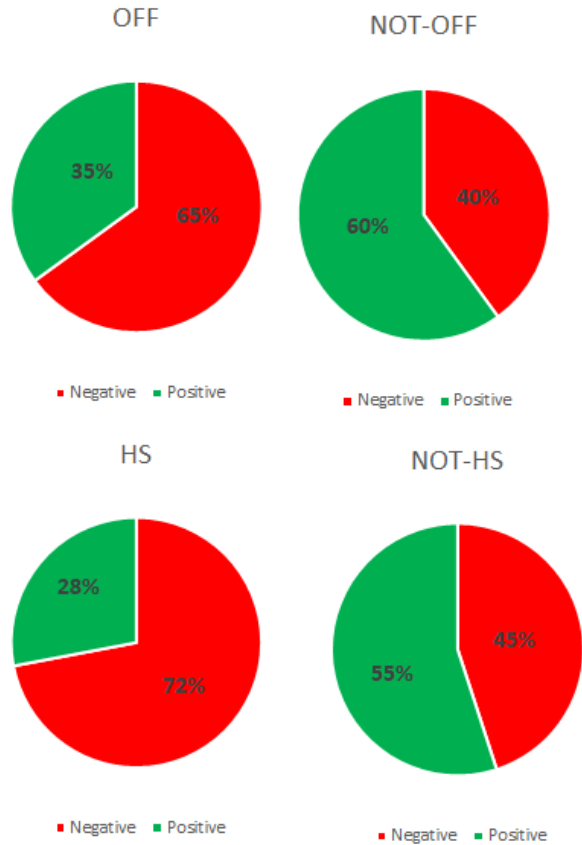


Figure 3: Distribution of Negative and Positive Tweets after applied AraNet on Shared-Task TRAIN Data

word pieces) and use the hidden states of these masked tokens from last layer for prediction. To understand the relationship between two sentences, the BERT also pre-trains with a binarized NSP task, which is also a type of self-supervised learning. For the sentence pairs (e.g., A-B) in pre-training examples, 50% of the time B is the actual next sentence that follows A in the corpus (positive class) and 50% of the time B is a random sentence from corpus (negative class). Google’s pre-trained BERT-Base Multilingual Cased model is trained on 104 languages (including Arabic) with 12 layers, 768 hidden units each, 12 attention heads. The model has 119,547 shared word pieces vocabulary, and was pre-trained on the entire Wikipedia for each language.

In our experiments, we train our classification models on BERT-Base Multilingual Cased model. For all of our fine-tuning BERT models, we use a maximum sequence size of 50 tokens and a batch size of 32. We add a ‘[CLS]’ token at the beginning of each input sequence and, then, feed the final hidden state of ‘[CLS]’ to a Softmax linear layer to get predication probabilities across classes. We set the learning rate to $2e - 6$ and train for 20 epochs. We save the checkpoint at the end of each epoch, report F1-score and accuracy of the best model, and use the best checkpoint to predict the labels of the TEST set. We fine-tune the BERT model under five settings. We describe each of these next.

Model	Dev				Test			
	OFF		HS		OFF		HS	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1
BERT	87.10	78.38	95.70	70.96	87.30	77.70	95.20	70.51
BERT-SENTI	87.40	78.84	95.50	68.01	87.45	80.51	93.15	61.57
BERT-EMO	88.30	80.39	95.40	68.54	–	–	–	–
BERT-EMO-AUG	89.60	82.31	93.90	62.52	89.35	82.85	–	–

Table 4: Offensive (OFF) and Hate Speech (HS) results on DEV and TEST datasets

Vanilla BERT: We fine-tune BERT-Base Multilingual Cased model on TRAIN set of offensive task and hate speech task respectively. We refer these two models to BERT. The offensive model obtains the best result with 8 epochs. As Table 4 shows, for *offensive language* classification, this model obtains 87.10% accuracy and 78.38 F_1 score on DEV set. We submit the TEST prediction of this model to the shared task and obtain 87.30% accuracy and 77.70 F_1 on the TEST set. The *hate speech* model obtains best result (accuracy = 95.7%, $F_1 = 70.96$) with 6 epochs.

BERT-SENTI We use a BERT model fine-tuned with on binary Arabic sentiment dataset as released by (Abdul-Mageed et al., 2020). We use this off-the-shelf (already trained) model to further fine-tune on offensive and hate speech tasks, respectively. We replace the Softmax linear layer for sentiment classification with a randomly initialized Softmax linear layer for each task. We refer to these two models as BERT-SENTI. We train the BERT-SENTI models on the TRAIN sets for offensive and hate speech tasks respectively. On F_1 score, BERT-SENTI is 0.3 better than vanilla BERT on the offensive task, but 2.95 lower (than vanilla BERT) on the hate speech task. We submit the TEST predictions of both tasks. The offensive model obtain 87.45% accuracy and 80.51 F_1 on TEST. The hate speech model acquire 93.15% accuracy and 61.57 F_1 on TEST.

BERT-EMO Similar to BERT-SENTI, we use a BERT model trained on 8-class Arabic emotion identification from (Abdul-Mageed et al., 2020) to fine-tune on the offensive and hate speech tasks, respectively. We refer to this setting as BERT-EMO. We train the models on the TRAIN sets for both offensive and hate speech tasks for 20 epochs. The *offensive* model obtains its best result (accuracy = 88.30%, $F_1 = 80.39$) with 11 epochs. The *hate speech* model acquires its best result (accuracy = 95.40%, $F_1 = 68.54$) also with 11 epochs. We do not submit an BERT-EMO on the hate speech task TEST set.

BERT-EMO-AUG Similar to BERT-EMO, we also fine-tune the emotion BERT model (BERT-EMO) with the augmented offensive dataset (AUG-TRAIN-OFF) and augmented hate speech dataset (AUG-TRAIN-HS). On the DEV set, the *offensive model* acquires its best result (accuracy = 89.60%, $F_1 = 82.31$) with 13 epochs. The best results for the *hate speech model* (accuracy = 93.90%, $F_1 = 62.52$) is obtained with 9 epochs. Our best offensive predi-

cation on TEST is BERT-EMO-AUG. It which achieves an accuracy of 89.35% and F_1 of 82.85. We do not submit an BERT-EMO-AUG on the hate speech task TEST set.

5. Conclusion

We described our submission to the offensive language detection in Arabic shared task. We offered a simple method to extend training data and demonstrated the utility of such augmented data empirically. We also deploy affective language models on the two sub-tasks of offensive language detection and hate speech identification. We show that fine-tuning such affective models is useful, especially in the case of offensive language detection. In the future, we will investigate other methods for improving our automatic offensive and hateful language acquisition methods. We also explore other machine learning methods on the tasks. For example, we plan to investigate the utility of semi-supervised methods as a vehicle of improving our models.

6. Bibliographic References

- Abdul-Mageed, M., Zhang, C., Elmadany, A., Rajendran, A., and Ungar, L. (2019). Dianet: Bert and hierarchical attention multi-task learning of fine-grained dialect. *arXiv preprint arXiv:1910.14243*.
- Abdul-Mageed, Muhammad, Z. C., Nagoudi, E. M. B., and Hashemi, A. (2020). Aranet: A deep learning toolkit for arabic social media. In *The 4th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT4), LREC*.
- Abozinadah, E. (2017). *Detecting abusive arabic language twitter accounts using a multidimensional analysis model*. Ph.D. thesis.
- Agrawal, S. and Awekar, A. (2018). Deep learning for detecting cyberbullying across multiple social media platforms. In *European Conference on Information Retrieval*, pages 141–153. Springer.
- Alakrot, A., Murray, L., and Nikolov, N. S. (2018). Towards accurate detection of offensive language in online communication in arabic. *Procedia computer science*, 142:315–320.
- 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)*, pages 69–76. IEEE.
- Badjatiya, P., Gupta, S., Gupta, M., and Varma, V. (2017). Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 759–760.

- Barberá, P. and Sood, G. (2015). Follow your ideology: Measuring media ideology on social networks. In *Annual Meeting of the European Political Science Association, Vienna, Austria*. Retrieved from <http://www.gsood.com/research/papers/mediabias.pdf>.
- Chadefaux, T. (2014). Early warning signals for war in the news. *Journal of Peace Research*, 51(1):5–18.
- Conover, M. D., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., and Flammini, A. (2011). Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media*.
- Dadvar, M., Trieschnigg, D., Ordelman, R., and de Jong, F. (2013). Improving cyberbullying detection with user context. In *European Conference on Information Retrieval*, pages 693–696. Springer.
- Darwish, K., Alexandrov, D., Nakov, P., and Mejova, Y. (2017). Seminar users in the arabic twitter sphere. In *International Conference on Social Informatics*, pages 91–108. Springer.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Fortuna, P., Ferreira, J., Pires, L., Routar, G., and Nunes, S. (2018). Merging datasets for aggressive text identification. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 128–139.
- Georgakopoulos, S. V., Tasoulis, S. K., Vrahatis, A. G., and Plagianakos, V. P. (2018). Convolutional neural networks for toxic comment classification. In *Proceedings of the 10th Hellenic Conference on Artificial Intelligence*, pages 1–6.
- Jay, T. and Janschewitz, K. (2008). The pragmatics of swearing. *Journal of Politeness Research. Language, Behaviour, Culture*, 4(2):267–288.
- Kumar, R., Ojha, A. K., Malmasi, S., and Zampieri, M. (2018). Benchmarking aggression identification in social media. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 1–11.
- Kwok, I. and Wang, Y. (2013). Locate the hate: Detecting tweets against blacks. In *Twenty-seventh AAAI conference on artificial intelligence*.
- Malmasi, S. and Zampieri, M. (2017). Detecting hate speech in social media. *arXiv preprint arXiv:1712.06427*.
- Modha, S., Majumder, P., and Mandl, T. (2018). Filtering aggression from the multilingual social media feed. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 199–207.
- Mubarak, H. and Darwish, K. (2019). Arabic offensive language classification on twitter. In *International Conference on Social Informatics*, pages 269–276. Springer.
- 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.
- Nobata, C., Tetreault, J., Thomas, A., Mehdad, Y., and Chang, Y. (2016). Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pages 145–153.
- Weber, I., Garimella, V. R. K., and Batayneh, A. (2013). Secular vs. islamist polarization in egypt on twitter. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining*, pages 290–297.
- Wiegand, M., Siegel, M., and Ruppenhofer, J. (2018). Overview of the germeval 2018 shared task on the identification of offensive language.
- Zampieri, M., Malmasi, S., Nakov, P., Rosenthal, S., Farra, N., and Kumar, R. (2019). Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). *arXiv preprint arXiv:1903.08983*.