

# Multilingual Offensive Language Identification for Low-resource Languages

THARINDU RANASINGHE, University of Wolverhampton, UK

MARCOS ZAMPIERI, Rochester Institute of Technology, USA

Offensive content is pervasive in social media and a reason for concern to companies and government organizations. Several studies have been recently published investigating methods to detect the various forms of such content (e.g. hate speech, cyberbullying, and cyberaggression). The clear majority of these studies deal with English partially because most annotated datasets available contain English data. In this paper, we take advantage of available English datasets by applying cross-lingual contextual word embeddings and transfer learning to make predictions in low-resource languages. We project predictions on comparable data in Arabic, Bengali, Danish, Greek, Hindi, Spanish, and Turkish. We report results of 0.8415 F1 macro for Bengali in TRAC-2 shared task [23], 0.8532 F1 macro for Danish and 0.8701 F1 macro for Greek in OffensEval 2020 [58], 0.8568 F1 macro for Hindi in HASOC 2019 shared task [27] and 0.7513 F1 macro for Spanish in in SemEval-2019 Task 5 (HatEval) [7] showing that our approach compares favourably to the best systems submitted to recent shared tasks on these three languages. Additionally, we report competitive performance on Arabic, and Turkish using the training and development sets of OffensEval 2020 shared task. The results for all languages confirm the robustness of cross-lingual contextual embeddings and transfer learning for this task.

CCS Concepts: • **Computing methodologies** → **Neural networks; Transfer learning; Supervised learning by classification.**

Additional Key Words and Phrases: offensive language identification, cross-lingual embeddings, low-resource languages

## ACM Reference Format:

Tharindu Ranasinghe and Marcos Zampieri. 2020. Multilingual Offensive Language Identification for Low-resource Languages. 1, 1 (May 2020), 13 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

Offensive posts on social media result in a number of undesired consequences to users. They have been investigated as triggers of suicide attempts and ideation, and mental health problems [5, 9]. One of the most common ways to cope with offensive content online is training systems capable of recognizing offensive messages or posts. Once recognized, such offensive content can be set aside for human moderation or deleted from the respective platform (e.g. Facebook, Twitter) preventing harm to users and controlling the spread of abusive behavior in social media.

There have been several recent studies published on this topic focusing on different aspects of offensiveness such as abuse [28], aggression [22, 23], cyber-bullying [42], and hate speech [25, 26, 45] to name a few. While there are a few studies published on languages such as Arabic [29] and Greek [35], most studies and datasets created thus far have focused on English. Data augmentation [15] and multilingual word embeddings [31] have been applied to take advantage of existing English datasets to improve the performance in systems dealing with languages other than

---

Authors' addresses: Tharindu Ranasinghe, University of Wolverhampton, Wulfruna Street, Wolverhampton, UK, WV1 1LY, T.D. RanasingheHettiarachchige@wlv.ac.uk; Marcos Zampieri, Rochester Institute of Technology, 92 Lomb Memorial Drive, Rochester, NY, USA, 14620, marcos.zampieri@rit.edu.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2020 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

English. To the best of our knowledge, however, state-of-the-art cross-lingual contextual embeddings such as XLM-R [11] have not yet been applied to offensive language identification. To address this gap, we evaluate the performance of cross-lingual contextual embeddings and transfer learning (TL) methods in projecting predictions from English to other languages. We show that our methods compare favourably to state-of-the-art approaches submitted to recent shared tasks on all datasets. The main contributions of this paper are the following:

- (1) We apply cross-lingual contextual word embeddings to offensive language identification. We take advantage of existing English data to project predictions in seven other languages: Arabic, Bengali, Danish, Greek, Hindi, Spanish, and Turkish.
- (2) We tackle both off-domain and off-task data for Bengali. We show that not only these methods can project predictions on different languages but also on different domains (e.g. Twitter vs. Facebook) and tasks (e.g. binary vs. three-way classification).
- (3) We provide important resources to the community: the code, and the English model will be freely available to everyone interested in working on low-resource languages using the same methodology.

Finally, it is worth noting that the seven languages other than English included in this study are all languages with millions of speakers. Their situation in terms of language resources is more favorable and not comparable to the situation of minority or endangered languages for which there are virtually no datasets available, often called truly low-resource languages [1] or extremely low-resource languages [48]. That said, we consider these seven languages as low-resourced in the context of offensive language identification due to the lack of large available offensive language datasets in these languages especially when compared to English. The methods presented in this paper can be used to truly low-resource languages addressing data scarcity.

## 2 RELATED WORK

There have been different types of abusive content addressed in recent studies including hate speech [26], aggression [22, 23], and cyberbullying [42]. A few annotation taxonomies, such as the one proposed by OLID [56] and replicated in other studies [43], try to take advantage of the similarities between these sub-tasks allowing us to consider multiple types of abusive language at once.

In terms of computational approaches, a variety of methods have been tested to identify abusive posts automatically including simple profanity lexicons, machine learning classifiers like Naive Bayes and SVMs, and deep learning representations and neural networks like convolutional neural networks (CNN), recurrent neural networks (RNN) [16], LSTM [4, 24] and transformers [40]. In this section we briefly summarise related work addressing the most common abusive language phenomena.

**Aggression identification:** The two editions of the TRAC shared task on Aggression Identification organized in 2018 and 2020 [22, 23] provided participants with annotated Facebook dataset containing posts and comments labeled as non-aggressive, covertly aggressive, and overtly aggressive. The best-performing systems in the 2018 edition competition used deep learning approaches based on CNN, RNN, and LSTM [4, 24]. In Section 5 we compare our results for Bengali with the best performing methods of TRAC 2020.

**Cyberbullying detection:** Cyberbullying is another popular topic with several studies published such as [55] who used sentiment analysis methods and topic modelling and [12] who addressed the problem using profiling-related features like the frequency of curse words in users' messages.

**Hate Speech:** Hate speech is by far the most studied phenomenon with several studies published on various languages [10, 14]. The recent HatEval [7] competition at SemEval 2019 addressed hate speech against women and migrants featuring datasets in English and Spanish.

**Offensive Language:** In addition to the popular OffensEval shared task, which introduced a taxonomy encompassing multiple types of abusive and offensive content discussed in more detail in Section 5.1, there have been several studies published on identifying offensive posts [35]. The GermEval [54] shared task provided participants with a dataset containing 8,500 annotated German tweets. Two sub-tasks were organized, the first sub-task was to discriminate between offensive and non-offensive tweets and the second sub-task was to discriminate between profanity, insult, and abuse.

**Other Languages:** While the clear majority of studies deal with English, there have been a number of studies dealing of other languages as well. Examples included Dutch [49, 50], German [44], Italian [32], and the languages studied in the present paper.

### 3 DATA

To carry out the experiments presented in this paper, we acquire datasets in English and other seven languages: Arabic, Bengali, Danish, Greek, Hindi, Spanish, and Turkish listed in Table 1.

Lang.	Instances	Source	Labels
Arabic	8,000	T	offensive, non-offensive
Bengali	4,000	Y	overtly aggressive, covertly aggressive, non aggressive
Danish	2,961	F, R	offensive, non-offensive
English	14,100	T	offensive, non-offensive
Greek	8,743	T	offensive, non-offensive
Hindi	8,000	T	hate offensive, non hate-offensive
Spanish	6,600	T	hateful, non-hateful
Turkish	31,756	T	offensive, non-offensive

Table 1. Instances (Inst.), source (S) and labels in all datasets. F stands for Facebook, R for Reddit, T for Twitter and Y for Youtube.

The Bengali, Hindi, and Spanish datasets have been used in shared tasks in 2019 and 2020 allowing us to compare the performance of our methods to other approaches. The Hindi dataset [27] was used in the HASOC 2019 shared task while the Spanish dataset [7] was used in SemEval-2019 Task 5 (HatEval). They both contain Twitter data and two labels. The Bengali dataset [8] was used in the TRAC-2 shared task [23] on aggression identification. It is different than the other three datasets in terms of domain (Youtube instead of Twitter) and set of labels (three classes instead of binary) allowing us to compare the performance of cross-lingual embeddings on off-domain data and off-task data.

The Arabic, Danish, Greek, and Turkish datasets are official datasets of the OffensEval 2020 competition hosted at SemEval 2020 (Task 12) [58]. We report the results, we obtained for Danish and Greek test sets. in Section 5.1. However, since Arabic and Turkish are morphologically rich, they require language specific segmentation approaches which we hope to conduct in the future. Therefore we only report results on the development set in Section 5.1 and a detailed explanation of future work on Arabic and Turkish is reported in Section 6.

Finally, as our English dataset, we choose the Offensive Language Identification Dataset (OLID) [56], used in the SemEval-2019 Task 6 (OffensEval) [57]. OLID is arguably one of the most popular offensive language datasets. It contains manually annotated tweets with the following three-level taxonomy and labels:

- A:** Offensive language identification - offensive vs. non-offensive;
- B:** Categorization of offensive language - targeted insult or thread vs. untargeted profanity;
- C:** Offensive language target identification - individual vs. group vs. other.

We chose OLID due to the flexibility provided by its hierarchical annotation model which considers multiple types of offensive content in a single taxonomy (e.g. targeted insults to a group are often *hate speech* whereas targeted insults to an individual are often *cyberbullying*). This allows us to map OLID level A (offensive vs. non-offensive) to labels in the other three datasets. OLID’s annotation model is intended to serve as a general-purpose model for multiple *abusive language detection sub-tasks* as described by Waseem et al. (2017) [53]. The transfer learning strategy used in this paper provides us with the opportunity to evaluate how closely the OLID labels relate to the classes in other datasets which were annotated using different guidelines and definitions (e.g. *aggression* and *hate speech*).

## 4 METHODOLOGY

Transformer models have been used successfully for various NLP tasks [13] such as text classification [18, 37], NER [39, 47], context similarity [17], language identification [19] etc. Most of the tasks were focused on English language due to the fact the most of the pre-trained transformer models were trained on English data. Even though, there were several multilingual models like BERT-m [13] there were many speculations about its ability to represent all the languages [34] and although BERT-m model showed some cross-lingual characteristics it has not been trained on crosslingual data [20]. The motivation behind this methodology was the recently released cross-lingual transformer models - XLM-R [11] which has been trained on 104 languages. The interesting fact about XLM-R is that it is very compatible in monolingual benchmarks while achieving best results in cross-lingual benchmarks at the same time [11]. The main idea of the methodology is that we train a classification model on a resource rich, typically English, using a cross-lingual transformer model and perform transfer learning on a less resource language.

There are two main parts of the methodology. Subsection 4.1 describes the classification architecture we used for all the languages. In Subsection 4.2 we describe the transfer learning strategies we used to utilise English offensive language data in predicting offense in less-resourced languages.

### 4.1 XLM-R for Text Classification

Similar to other transformer architectures XLM-R transformer architecture can also be used for text classification tasks [11]. XLM-R-large model contains approximately 125M parameters with 12-layers, 768 hidden-states, 3072 feed-forward hidden-states and 8-heads [11]. It takes an input of a sequence of no more than 512 tokens and outputs the representation of the sequence. The first token of the sequence is always [CLS] which contains the special classification embedding [46].

For text classification tasks, XLM-R takes the final hidden state  $h$  of the first token [CLS] as the representation of the whole sequence. A simple softmax classifier is added to the top of XLM-R to predict the probability of label  $c$ : as shown in Equation 1 where  $W$  is the task-specific parameter matrix [18, 37].

$$p(c|h) = \text{softmax}(Wh) \quad (1)$$

We fine-tune all the parameters from XLM-R as well as  $W$  jointly by maximising the log-probability of the correct label. The architecture diagram of the classification is shown in Figure 1. We specially used the XLM-R large model.

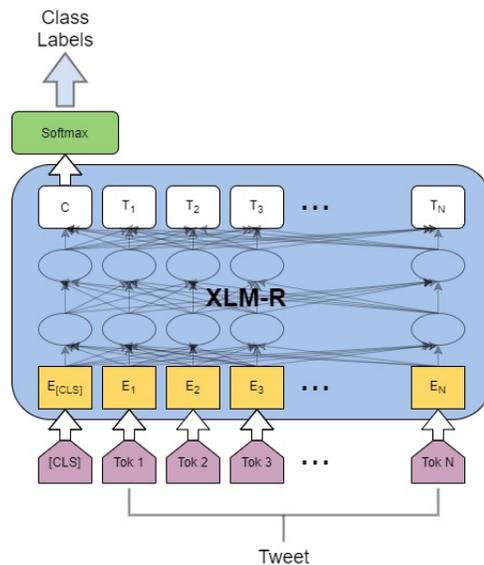


Fig. 1. Text Classification Architecture [38]

#### 4.2 Transfer-learning strategies

When we adopt XLM-R for multilingual offensive language identification, we perform transfer learning in two different ways.

*Inter-language transfer learning.* We first trained the XLM-R classification model on first level of English offensive language identification dataset (OLID) [56]. Then we save the weights of the XLM-R model as well as the softmax layer. We use this saved weights from English to initialise the weights for a new language. To explore this transfer learning aspect we experimented on Hindi language which was released for HASOC 2019 shared task [27], on Spanish data released for Hateval 2019 [7] and on Arabic, Danish, Greek and Turkish data released for OffensEval 2020 [58].

*Inter-task and inter-language transfer learning.* Similar to inter-language transfer learning strategy, we first trained the XLM-R classification model on the first level of English offensive language identification dataset (OLID) [56]. Then we only save the weights of the XLM-R model and use the saved weights to initialise the weights for a new language. We did not use the weights of the last softmax layer since we wanted to test this strategy on different data that has a different number of offensive classes to predict. Since the last softmax layer reflects the number of classes, it is not possible to transfer the weights of the softmax layer when the number of the classes are different in the classification task. We explored this transfer learning aspect with Bengali dataset released with TRAC - 2 shared task [23]. As described in the Section 3 the classifier should make a 3-way classification in between 'Overtly Aggressive', 'Covertly Aggressive' and 'Non Aggressive' text data.

We used a Nvidia Tesla K80 GPU to train the models. We divided the data set into a training set and a validation set using 0.8:0.2 split on the dataset. We mainly fine tuned the learning rate and number of epochs of the classification model manually to obtain the best results for the validation set. We obtained  $1e^{-5}$  as the best value for learning rate and 3 as the best value for number of epochs for all the languages. The other configurations of the transformer model were set to a constant value over all the languages in order to ensure consistency between the languages. This also provides a good starting configuration for researchers who intend to use this technique on a new language. We used a batch-size of eight, Adam optimiser [21] and a linear learning rate warm-up over 10% of the training data. The models were trained using only training data. We performed early stopping if the evaluation loss did not improve over ten evaluation rounds.

Training for English language took around 1 hour while training for other languages took around 30 minutes.

## 5 RESULTS AND EVALUATION

We evaluate the results obtained by all models using the test set provided by the organizers of each competition. We compared our results to the best systems in TRAC-2 for Bengali, HASOC for Hindi, HatEval for Spanish in terms of weighted and macro F1 score according to the metrics reported by the task organizers - TRAC-2 reported only macro F1, HatEval reported only weighted F1 and HASOC reported both macro F1 and weighted F1. Finally, we quantify the improvement of the transfer learning strategy in the performance of both BERT and XLM-R. *TL* indicates that the model used the transfer learning strategy described in Subsection 4.2.

Model	Non Hate Offensive			Hate Offensive			Weighted Average			F1 Macro
	P	R	F1	P	R	F1	P	R	F1	
<i>XLM-R (TL)</i>	0.84	0.87	0.88	0.84	0.86	0.85	0.86	0.85	0.86	0.86
<i>BERT-m (TL)</i>	0.80	0.83	0.84	0.80	0.82	0.81	0.82	0.81	0.82	0.82
<i>Bashar and Nayak [6]</i>	NR	NR	NR	NR	NR	NR	NR	NR	0.82	0.81
<i>XLM-R</i>	0.84	0.96	0.90	0.86	0.61	0.71	0.85	0.85	0.81	0.80
<i>BERT-m</i>	0.77	0.99	0.86	0.94	0.33	0.49	0.82	0.78	0.80	0.80

Table 2. Results for offensive language detection in Hindi. For each model, Precision (P), Recall (R), and F1 are reported on all classes, and weighted averages. Macro-F1 is also listed. NR indicates that the research did not report that value. Bashar and Nayak [6] is the best result submitted to the competition.

Model	Non Hateful			Hateful			Weighted Average			F1 Macro
	P	R	F1	P	R	F1	P	R	F1	
<i>XLM-R (TL)</i>	0.78	0.76	0.77	0.76	0.78	0.73	0.76	0.76	0.76	0.75
<i>BERT-m (TL)</i>	0.76	0.74	0.75	0.74	0.76	0.71	0.74	0.74	0.74	0.73
<i>Vega et al. [51]</i>	NR	NR	NR	NR	NR	NR	NR	NR	0.73	NR
<i>Pérez and Luque [33]</i>	NR	NR	NR	NR	NR	NR	NR	NR	0.73	NR
<i>XLM-R</i>	0.77	0.74	0.74	0.73	0.76	0.70	0.73	0.73	0.73	0.72
<i>BERT-m</i>	0.75	0.72	0.73	0.71	0.74	0.68	0.72	0.72	0.72	0.71

Table 3. Results for offensive language detection in Spanish. For each model, Precision (P), Recall (R), and F1 are reported on all classes, and weighted averages. Macro-F1 is also listed. NR indicates that the research did not report that value. Vega et al. [51] and Pérez and Luque [33] are the best results submitted to the competition.

Model	Non Aggressive			Overtly Aggressive			Covertly Aggressive			Weighted Average			F1 Macro
	P	R	F1	P	R	F1	P	R	F1	P	R	F1	
<i>XLM-R (TL)</i>	0.96	0.94	0.95	0.73	0.81	0.77	0.72	0.80	0.76	0.86	0.85	0.84	0.84
<i>Risch and Krestel [41]</i>	NR	NR	NR	NR	NR	NR	NR	NR	NR	NR	NR	0.82	0.82
<i>BERT-m (TL)</i>	0.94	0.92	0.93	0.71	0.79	0.75	0.70	0.78	0.68	0.80	0.83	0.82	0.82
<i>XLM-R</i>	0.91	0.90	0.91	0.70	0.78	0.74	0.69	0.77	0.67	0.79	0.83	0.82	0.81
<i>BERT-m</i>	0.90	0.89	0.90	0.69	0.77	0.73	0.68	0.76	0.66	0.78	0.81	0.81	0.81

Table 4. Results for offensive language detection in Bengali. For each model, Precision (P), Recall (R), and F1 are reported on all classes, and weighted averages. Macro-F1 is also listed. NR indicates that the research did not report that value. Risch and Krestel [41] is the best result submitted to the competition.

For Hindi, as presentend in Table 2, transfer learning with XLM-R cross lingual embeddings provided the best results achieving 0.86 for both weighted and macro F1 score. In HASOC 2019 [27], the best model by Bashar and Nayak [6] scored 0.81 Macro F1 and 0.82 Weighted F1 using convolutional neural networks.

For Spanish transfer learning with XLM-R cross lingual embeddings also provided the best results achieving 0.75 and 0.76 macro and weighted F1 score respectively. The best two models in HatEval [7] for Spanish scored 0.73 macro F1 score. Both models applied SVM classifiers trained on a variety of features like character and word n-grams, POS tags, offensive word lexica, and embeddings.

The results on Bengali shown in Table 4 deserve special attention because the Bengali data is off-domain with respect to the English data (Facebook instead of Twitter) and it contains three labels (covertly aggressive, overtly aggressive, and not aggressive) instead of two in the English dataset (offensive and non-offensive). *TL* indicates that the model used the inter-task, inter-domain, and inter-language transfer learning strategy described in Subsection 4.2. Similar to the Hindi and Spanish, for Bengali also transfer learning with XLM-R cross lingual embeddings provided the best results achieving 0.75 and 0.76 macro and weighted F1 respectively thus outperforming the other models by a significant margin. The best model in the TRAC-2 shared task [23] scored 0.82 weighted F1 score in Bengali using a BERT-based system.

We look closer to the test set predictions by XLM-R (TL) for Bengali in Figure 2. We observe that the performance for the non-aggressive class is substantially better than the performance for the overtly aggressive and covertly aggressive classes following a trend observed by the TRAC-2 participants including Risch and Krestel [41].

Finally, it is clear that in all the experimental settings, the cross-lingual embedding models fine-tuned with transfer learning, outperforms the best system available for the three languages. Furthermore, the results show that the cross-lingual nature of the XLM-R model provided a boost over the multilingual BERT model in all languages tested.

## 5.1 OffensEval 2020

We also experimented our methods with the recently released OffensEval 2020 [58] datasets in four languages: Arabic, Danish, Greek and Turkish. In Danish and Turkish we evaluate our results on the test sets comparing them to the best systems in the competition. As shown in Table 5 and Table 6 our methodology out performs the best systems submitted in each of these languages. As we mentioned earlier, remaining languages of OffensEval 2020 [58]: Arabic and Turkish are morphologically rich languages and they require language specific segmentation techniques which we hope to explore in future research. That said, in order to show our transfer learning strategy works for a wider range of languages, we presents the results we obtained for development set in the Tables 7 and 8. Although not fully comparable, the high results obtained on the development set provides us an indication that the proposed approach is



Fig. 2. Heat map of the Bengali test set predictions by XLM-R (TL).

likely to obtain high results on the Arabic and Turkish test sets as well. We aim to perform this evaluation with the different morphological segmentation techniques in both Arabic and Turkish.

Finally, in all of the OffensEval 2020 languages, XLM-R with the transfer learning strategy performs best from the experiments. It is clear that transfer learning approach boosts the performance of both BERT-m and XLM-R models through out all the languages. Furthermore, it is clear that there is a significant advantage in using cross lingual transformers with transfer learning rather than training transformer models from scratch.

Model	Not Offensive			Offensive			Weighted Average			F1 Macro
	P	R	F1	P	R	F1	P	R	F1	
<i>XLM-R (TL)</i>	0.85	0.88	0.89	0.69	0.75	0.72	0.87	0.87	0.87	0.83
<i>BERT-m (TL)</i>	0.89	0.87	0.88	0.67	0.73	0.70	0.85	0.86	0.86	0.82
<i>Pamies et al. [30]</i>	NR	NR	NR	NR	NR	NR	NR	NR	NR	0.81
<i>XLM-R</i>	0.87	0.85	0.86	0.66	0.72	0.70	0.85	0.85	0.85	0.81
<i>BERT-m</i>	0.86	0.84	0.85	0.65	0.71	0.69	0.84	0.84	0.84	0.80

Table 5. Results for offensive language detection in Danish. For each model, Precision (P), Recall (R), and F1 are reported on all classes, and weighted averages. Macro-F1 is also listed. NR indicates that the research did not report that value. Pamies et al. [30] is the best result submitted to the competition.

Model	Not Offensive			Offensive			Weighted Average			F1 Macro
	P	R	F1	P	R	F1	P	R	F1	
<i>XLM-R (TL)</i>	0.95	0.93	0.94	0.72	0.80	0.76	0.91	0.91	0.91	0.87
<i>BERT-m (TL)</i>	0.94	0.92	0.93	0.71	0.79	0.75	0.90	0.90	0.90	0.86
<i>Ahn et al. [2]</i>	NR	NR	NR	NR	NR	NR	NR	NR	NR	0.85
<i>XLM-R</i>	0.91	0.89	0.90	0.68	0.75	0.71	0.86	0.86	0.87	0.83
<i>BERT-m</i>	0.90	0.88	0.89	0.67	0.74	0.70	0.85	0.85	0.85	0.81

Table 6. Results for offensive language detection in Greek. For each model, Precision (P), Recall (R), and F1 are reported on all classes, and weighted averages. Macro-F1 is also listed. NR indicates that the research did not report that value. Ahn et al. [2] is the best result submitted to the competition.

Model	Not Offensive			Offensive			Weighted Average			F1 Macro
	P	R	F1	P	R	F1	P	R	F1	
<i>XLM-R (TL)</i>	0.96	0.94	0.95	0.73	0.81	0.77	0.92	0.92	0.92	0.88
<i>BERT-m (TL)</i>	0.95	0.92	0.93	0.72	0.78	0.76	0.90	0.90	0.90	0.86
<i>XLM-R</i>	0.93	0.90	0.91	0.70	0.76	0.74	0.89	0.89	0.89	0.85
<i>BERT-m</i>	0.91	0.89	0.90	0.70	0.75	0.49	0.87	0.87	0.88	0.84

Table 7. Results for offensive language detection in Arabic. For each model, Precision (P), Recall (R), and F1 are reported on all classes, and weighted averages. Macro-F1 is also listed. NR indicates that the research did not report that value. Alami et al. [3] is the best result submitted to the competition.

Model	Not Offensive			Offensive			Weighted Average			F1 Macro
	P	R	F1	P	R	F1	P	R	F1	
<i>XLM-R (TL)</i>	0.83	0.81	0.82	0.70	0.78	0.74	0.89	0.89	0.89	0.85
<i>BERT-m (TL)</i>	0.81	0.79	0.80	0.69	0.77	0.72	0.87	0.88	0.88	0.84
<i>XLM-R</i>	0.79	0.77	0.78	0.67	0.75	0.70	0.85	0.85	0.86	0.80
<i>BERT-m</i>	0.77	0.75	0.77	0.66	0.74	0.69	0.84	0.84	0.85	0.79

Table 8. Results for offensive language detection in Turkish. For each model, Precision (P), Recall (R), and F1 are reported on all classes, and weighted averages. Macro-F1 is also listed. NR indicates that the research did not report that value. Wang et al. [52] is the best result submitted to the competition.

## 5.2 Progress Tests

The biggest challenge in developing deep learning models for offensive language identification in low resources languages is finding suitable annotated data. Therefore, we investigate whether our transfer learning approach can thrive in even less resource environments. In order to analyze this, we conduct the experiments for 0 (unsupervised), 100, 200, 300 and up to 1,000 training instances. We performed this in five languages: Hindi, Spanish, Danish, Greek and Bengali. We compare the results with and without transfer learning for both BERT-m and XLM-R. We report the Macro-F1 score of the predictions and gold labels in the test set against the number of training instances in Figure 3.

The transfer learning strategy significantly impacts the results. For Hindi, with only 100 training instances, training XLM-R from scratch achieves only 0.25 Macro F1 between the predictions and gold labels of the test set. However, with only 100 training instances, training XLM-R using the transfer learning strategy achieves 0.61 Macro F1 in Hindi. When the number of training instances grows, the results from the XLM-R models trained with the transfer learning strategy and the results from the XLM-R models trained from scratch converge. Similar pattern can be observed with BERT-m model too. This can be noticed in all the languages we experimented. However, in Bengali increasing number of training instances from 0 to 100 dropped the transfer learning results a bit, most probably because it takes certain number of examples to train the softmax layer we added.

It is interesting to see that, in all the languages, transfer learning strategy with zero number of training instances provided a good Macro F1 score. Therefore, it is clear that the cross lingual nature of these transformer models can identify offensive content in an unsupervised environment too, once we do transfer learning from English.

Our results confirm that XLM-R with the transfer learning strategy can be hugely beneficial to low-resource languages in offensive language identification where annotated training instances are scarce.

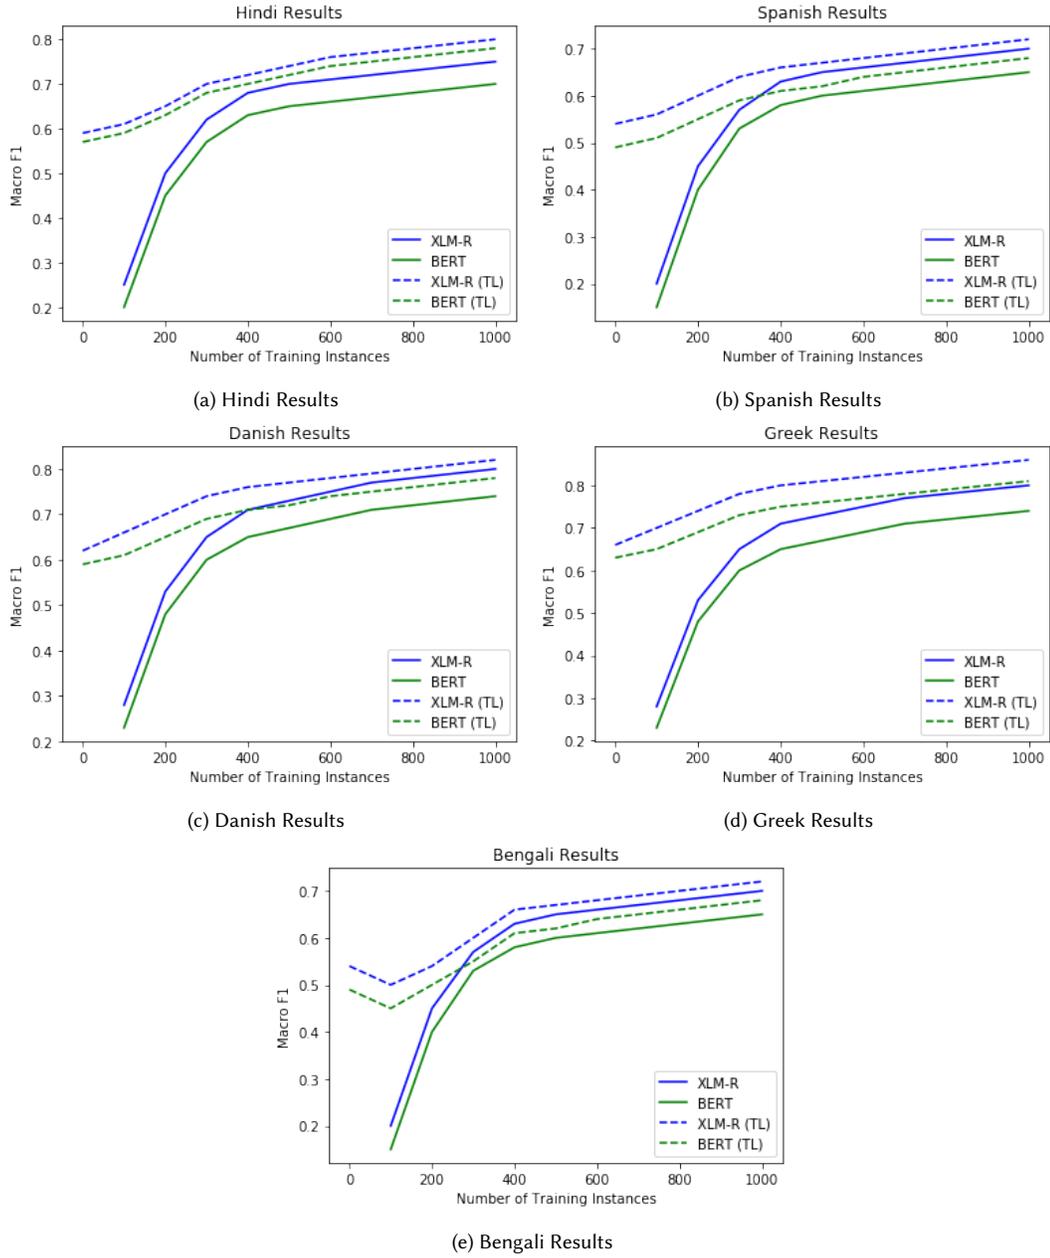


Fig. 3. Transfer learning impact in Offensive Language Identification. **XLM-R** and **BERT-m** indicates that the model was trained from scratch while **XLM-R (TL)** and **BERT (TL)** indicates that models followed the transfer learning strategy.

## 6 CONCLUSION AND FUTURE WORK

This paper is the first study to apply cross-lingual contextual word embeddings in offensive language identification projecting predictions from English to other languages. We developed systems based on different transformer architectures and evaluated their performance on benchmark datasets on Bengali [23], Hindi [27], and Spanish [7]. Finally, we compared them with the performance obtained by the systems that participated in shared tasks and we have showed that our best models outperformed the best systems in these competitions. Furthermore, our methods obtained promising results on the Arabic, Danish, Greek, and Turkish training and development sets released as part of the recent OffensEval 2020 [58].

We have showed that XLM-R with transfer learning outperforms all other methods we tested as well as the best results obtained by participants of the three competitions. Furthermore, the results from Bengali show that it is possible to achieve high performance using transfer learning on off-domain and off-task data when the labels do not have a direct correspondence in the projected dataset (two in English and three in Bengali). This opens exciting new avenues for future research considering the multitude of phenomena (e.g. hate speech, abuse, cyberbullying), annotation schemes, and guidelines used in offensive language datasets.

We are keen to cover more languages for offensive language identification using this strategy. Furthermore, when we are experimenting with morphologically rich languages like Arabic and Turkish we are interested to see whether language specific preprocessing like segmentation would improve the results. Finally, we would also like to apply our models to more low-resource languages creating important resources to cope with the problem of offensive language in social media.

We are currently experimenting with the recently finished HASOC 2020 shared task on identifying offensive YouTube comments in Code-mixed Malayalam. We are keen to experiment our transfer learning strategy in this task since the dataset is off domain and it contains code-mixed data. To the best of our knowledge, this is the first time that such methods are being tested on code-mixed data. Initial experiments [36] show that XLM-R with transfer learning strategy provides 0.85 and 0.73 weighted F1 and macro F1 scores respectively while training XLM-R from scratch gets only 0.83 and 0.71 weighted F1 and macro F1 scores respectively. These results suggest that our transfer learning approach is very likely to perform well with code-mixed data compared to other state-of-the-art methods.

## ACKNOWLEDGMENTS

We would like to thank the shared task organizers for making the datasets used in this paper available. We further thank the anonymous reviewers who provided us with constructive and insightful feedback to improve the quality of this paper.

## REFERENCES

- [1] Željko Agić, Anders Johannsen, Barbara Plank, Héctor Martínez Alonso, Natalie Schluter, and Anders Søgaard. 2016. Multilingual projection for parsing truly low-resource languages. *Transactions of the Association for Computational Linguistics* 4 (2016), 301–312.
- [2] Hwijee Ahn, Jimin Sun, Chan Young Park, and Jungyun Seo. 2020. NLPDove at SemEval-2020 Task 12: Improving Offensive Language Detection with Cross-lingual Transfer. In *Proceedings of SemEval*.
- [3] Hamza Alami, Said Ouatik El Alaoui, Abdessamad Benlahbib, and Nouredine En-nahnahi. 2020. LISAC FSDM-USMBA Team at SemEval 2020 Task 12: Overcoming AraBERT’s pretrain-finetune discrepancy for Arabic offensive language identification. In *Proceedings of SemEval*.
- [4] Segun Taofeek Aroyehun and Alexander Gelbukh. 2018. Aggression Detection in Social Media: Using Deep Neural Networks, Data Augmentation, and Pseudo Labeling. In *Proceedings of TRAC*.
- [5] Rienke Bannink, Suzanne Broeren, Petra M van de Looij-Jansen, Frouwkje G de Waart, and Hein Raat. 2014. Cyber and Traditional Bullying Victimization as a Risk Factor for Mental Health Problems and Suicidal Ideation in Adolescents. *PLoS one* 9, 4 (2014).

- [6] Md Abul Bashar and Richi Nayak. 2019. QutNocturnal@ HASOC'19: CNN for hate speech and offensive content identification in Hindi language. In *Proceedings of FIRE*.
- [7] Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. SemEval-2019 Task 5: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter. In *Proceedings of SemEval*.
- [8] Shiladitya Bhattacharya, Siddharth Singh, Ritesh Kumar, Akanksha Bansal, Akash Bhagat, Yogesh Dawer, Bornini Lahiri, and Atul Kr. Ojha. 2020. Developing a Multilingual Annotated Corpus of Misogyny and Aggression. In *Proceedings of TRAC*.
- [9] Rina A Bonanno and Shelley Hymel. 2013. Cyber bullying and internalizing difficulties: Above and beyond the impact of traditional forms of bullying. *Journal of youth and adolescence* 42, 5 (2013), 685–697.
- [10] Pete Burnap and Matthew L Williams. 2015. Cyber hate speech on twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy & Internet* 7, 2 (2015), 223–242.
- [11] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of ACL*.
- [12] Maral Dadvar, Dolf Trieschnigg, Roeland Ordelman, and Franciska de Jong. 2013. Improving cyberbullying detection with user context. In *Advances in Information Retrieval*. Springer, 693–696.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL*.
- [14] Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. 2015. Hate speech detection with comment embeddings. In *Proceedings of WWW*.
- [15] Erfan Ghadery and Marie-Francine Moens. 2020. LIIR at SemEval-2020 Task 12: A Cross-Lingual Augmentation Approach for Multilingual Offensive Language Identification. In *Proceedings of SemEval*.
- [16] Hansi Hettiarachchi and Tharindu Ranasinghe. 2019. Emoji Powered Capsule Network to Detect Type and Target of Offensive Posts in Social Media. In *Proceedings of RANLP*.
- [17] Hansi Hettiarachchi and Tharindu Ranasinghe. 2020. BRUMS at SemEval-2020 Task 3: Contextualised Embeddings for Predicting the (Graded) Effect of Context in Word Similarity. In *Proceedings of SemEval*.
- [18] Hansi Hettiarachchi and Tharindu Ranasinghe. 2020. InfoMiner at WNUT-2020 Task 2: Transformer-based Covid-19 Informative Tweet Extraction. In *Proceedings of W-NUT*.
- [19] Tommi Jauhiainen, Tharindu Ranasinghe, and Marcos Zampieri. 2021. Comparing Approaches to Dravidian Language Identification. In *Proceedings of VarDial*.
- [20] K Karthikeyan, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. Cross-lingual ability of multilingual BERT: An empirical study. In *Proceedings of ICLR*.
- [21] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*.
- [22] Ritesh Kumar, Atul Kr. Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking Aggression Identification in Social Media. In *Proceedings of TRAC*.
- [23] Ritesh Kumar, Atul Kr. Ojha, Shervin Malmasi, and Marcos Zampieri. 2020. Evaluating Aggression Identification in Social Media. In *Proceedings of TRAC*.
- [24] Prasenjit Majumder, Thomas Mandl, et al. 2018. Filtering Aggression from the Multilingual Social Media Feed. In *Proceedings of TRAC*.
- [25] Shervin Malmasi and Marcos Zampieri. 2017. Detecting Hate Speech in Social Media. In *Proceedings of RANLP*.
- [26] Shervin Malmasi and Marcos Zampieri. 2018. Challenges in Discriminating Profanity from Hate Speech. *Journal of Experimental & Theoretical Artificial Intelligence* (2018).
- [27] Thomas Mandl, Sandip Modha, Prasenjit Majumder, Daksh Patel, Mohana Dave, Chintak Mandlia, and Aditya Patel. 2019. Overview of the HASOC Track at FIRE 2019: Hate Speech and Offensive Content Identification in Indo-European Languages. In *Proceedings of FIRE*.
- [28] Hamdy Mubarak, Darwish Kareem, and Magdy Walid. 2017. Abusive Language Detection on Arabic Social Media. In *Proceedings of ALW*.
- [29] Hamdy Mubarak, Ammar Rashed, Kareem Darwish, Younes Samih, and Ahmed Abdelali. 2020. Arabic Offensive Language on Twitter: Analysis and Experiments. *arXiv preprint arXiv:2004.02192*.
- [30] Marc Pamies, Emily Ohman, Kaisla Kajava, and Jorg Tiedemann. 2020. LT@Helsinki at SemEval-2020 Task 12: Multilingual or language-specific BERT?. In *Proceedings of SemEval*.
- [31] Endang Wahyu Pamungkas and Viviana Patti. 2019. Cross-domain and Cross-lingual Abusive Language Detection: A Hybrid Approach with Deep Learning and a Multilingual Lexicon. In *Proceedings ACL:SRW*.
- [32] Serena Pelosi, Alessandro Maisto, Pierluigi Vitale, and Simonetta Vietri. 2017. Mining Offensive Language on Social Media. In *Proceedings of CLiC-it*.
- [33] Juan Manuel Pérez and Franco M Luque. 2019. Atalaya at SemEval 2019 task 5: Robust embeddings for tweet classification. In *Proceedings of SemEval*.
- [34] Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How Multilingual is Multilingual BERT?. In *Proceedings of ACL*.
- [35] Zeses Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. Offensive Language Identification in Greek. In *Proceedings of LREC*.
- [36] Tharindu Ranasinghe, Sarthak Gupte, Marcos Zampieri, and Ifeoma Nwogu. 2020. WLV-RIT at HASOC-Dravidian-CodeMix-FIRE2020: Offensive Language Identification in Code-switched YouTube Comments. In *Proceedings of FIRE*.
- [37] Tharindu Ranasinghe and Hansi Hettiarachchi. 2020. BRUMS at SemEval-2020 Task 12: Transformer Based Multilingual Offensive Language Identification in Social Media. In *Proceedings of SemEval*.

- [38] Tharindu Ranasinghe and Marcos Zampieri. 2020. Multilingual Offensive Language Identification with Cross-lingual Embeddings. In *Proceedings of EMNLP*.
- [39] Tharindu Ranasinghe and Marcos Zampieri. 2021. MUDES: Multilingual Detection of Offensive Spans. In *Proceedings of NAACL*.
- [40] Tharindu Ranasinghe, Marcos Zampieri, and Hansi Hettiarachchi. 2019. BRUMS at HASOC 2019: Deep Learning Models for Multilingual Hate Speech and Offensive Language Identification. In *Proceedings of FIRE*.
- [41] Julian Risch and Ralf Krestel. 2020. Bagging BERT Models for Robust Aggression Identification. In *Proceedings of TRAC*.
- [42] Hugo Rosa, N Pereira, Ricardo Ribeiro, Paula Costa Ferreira, Joao Paulo Carvalho, S Oliveira, Luísa Coheur, Paula Paulino, AM Veiga Simão, and Isabel Trancoso. 2019. Automatic cyberbullying detection: A systematic review. *Computers in Human Behavior* 93 (2019), 333–345.
- [43] Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2020. A Large-Scale Semi-Supervised Dataset for Offensive Language Identification. *arXiv preprint arXiv:2004.14454* (2020).
- [44] Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurovsky, and Michael Wojatzki. 2016. Measuring the Reliability of Hate Speech Annotations: The Case of the European Refugee Crisis. In *Proceedings of NLP4CMC*.
- [45] Paul Röttger, Bertram Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2020. HateCheck: Functional Tests for Hate Speech Detection Models. *arXiv preprint arXiv:2012.15606* (2020).
- [46] Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to Fine-Tune BERT for Text Classification?. In *Chinese Computational Linguistics*.
- [47] Ehsan Taher, Seyed Abbas Hoseini, and Mehrnosh Shamsfard. 2019. Beheshti-NER: Persian named entity recognition Using BERT. In *Proceedings of NSURL:ICNLSP*.
- [48] Allahsera Auguste Tapo, Bakary Coulibaly, Sébastien Diarra, Christopher Homan, Julia Kreutzer, Sarah Luger, Arthur Nagashima, Marcos Zampieri, and Michael Leventhal. 2020. Neural Machine Translation for Extremely Low-Resource African Languages: A Case Study on Bambara. In *Proceedings of LowResMT*.
- [49] Stéphan Tulkens, Lisa Hilde, Elise Lodewyckx, Ben Verhoeven, and Walter Daelemans. 2016. A Dictionary-based Approach to Racism Detection in Dutch Social Media. In *Proceedings of TA-COS*.
- [50] Cynthia Van Hee, Els Lefever, Ben Verhoeven, Julie Mennes, Bart Desmet, Guy De Pauw, Walter Daelemans, and Veronique Hoste. 2015. Automatic Detection and Prevention of Cyberbullying. In *Proceedings of HUSO*.
- [51] Luis Enrique Argota Vega, Jorge Carlos Reyes-Magaña, Helena Gómez-Adorno, and Gemma Bel-Enguix. 2019. MineríaUNAM at SemEval-2019 Task 5: Detecting Hate Speech in Twitter using Multiple Features in a Combinatorial Framework. In *Proceedings of SemEval*.
- [52] Shuohuan Wang, Jiaxiang Liu, Xuan Ouyang, and Yu Sun. 2020. Galileo at SemEval-2020 Task 12: Multi-lingual Learning for Offensive Language Identification using Pre-trained Language Models. In *Proceedings of SemEval*.
- [53] Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. Understanding Abuse: A Typology of Abusive Language Detection Subtasks. In *Proceedings of ALW*.
- [54] Michael Wiegand, Melanie Siegel, and Josef Ruppenhofer. 2018. Overview of the GermEval 2018 Shared Task on the Identification of Offensive Language. In *Proceedings of GermEval*.
- [55] Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012. Learning from bullying traces in social media. In *Proceedings of NAACL-HLT*.
- [56] Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Predicting the Type and Target of Offensive Posts in Social Media. In *Proceedings of NAACL*.
- [57] Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval). In *Proceedings of SemEval*.
- [58] Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. SemEval-2020 Task 12: Multilingual Offensive Language Identification in Social Media (OffensEval 2020). In *Proceedings of SemEval*.