

Team Rouges at SemEval-2020 Task 12: Cross-lingual Inductive Transfer to Detect Offensive Language

Tanvi Dadu * and Kartikey Pant *

Netaji Subhas Institute of Technology, New Delhi
International Institute of Information Technology, Hyderabad
tanvid.co.16@nsit.net.in
kartikey.pant@research.iiit.ac.in

Abstract

With the growing use of social media and its availability, many instances of the use of offensive language have been observed across multiple languages and domains. This phenomenon has given rise to the growing need to detect the offensive language used in social media cross-lingually. In OffensEval 2020, the organizers have released the *multilingual Offensive Language Identification Dataset* (mOLID), which contains tweets in five different languages, to detect offensive language. In this work, we introduce a cross-lingual inductive approach to identify the offensive language in tweets using the contextual word embedding *XLM-RoBERTa* (XLM-R). We show that our model performs competitively on all five languages, obtaining the fourth position in the English task with an F1-score of 0.919 and eighth position in the Turkish task with an F1-score of 0.781. Further experimentation proves that our model works competitively in a zero-shot learning environment, and is extensible to other languages.

1 Introduction

The prevalence of social media has made public commentary a critical aspect in shaping public opinion. Although freedom of speech is often advocated, offensive language in social media is unacceptable. Nevertheless, social media platforms and online communities are laden with offensive comments. This phenomenon results in the need for computationally identifying offense, aggression, and hate-speech in user-generated content in multiple languages.

This paper addresses the challenge put forward in the Multilingual Offensive Language Identification in Social Media shared task-organized at SemEval 2020 (Zampieri et al., 2020). The theme of the problem is to identify the offensive language in tweets in Arabic, Danish, English, Greek, and Turkish. This shared task is further divided into three sub-tasks. The first task consists of the identification of offensive tweets in a multilingual setting, whereas the other two tasks consist of the categorization of offensive tweets and identification of targets in English.

Transfer learning is a methodology to utilize the knowledge acquired from one or more tasks to solve other related tasks. It is ubiquitous in the domain of Natural Language Processing. It can be classified into multiple types, like transductive transfer learning and inductive transfer learning. The former is used when the tasks are the same, but the corresponding domains are different, such as in the case of the cross-lingual learning for similar tasks. On the other hand, the latter is used when the tasks are different, but the domains are similar such as in the case of finetuning of pretrained contextual word embeddings.

Cross-lingual inductive learning has been used for many downstream tasks like multi-lingual variants of question answering, text classification, and text generation (Artetxe and Schwenk, 2018; Lample and Conneau, 2019; Conneau et al., 2019). The rationale behind this is that the language with limited resources benefits from joint training over many languages. It also helps in performing zero-shot learning and handling of code-switched text, which otherwise are difficult to tackle.

In this work, we propose the use of cross-lingual inductive learning to detect the offensive language in the given five languages. We use pretrained *XLM-R* (Conneau et al., 2019) cross-lingual embeddings

* Both authors contributed equally to the work.

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

and train a single cross-lingual model for all five languages in the multilingual Offensive Language Identification Dataset.

2 Related Work

Xu et al. (2012) proposed the task of identifying bullying in social media using NLP. They presented benchmarking results for text classification among different NLP tasks using off-the-shelf solutions. Further, Nobata et al. (2016) proposed a machine learning-based method incorporating linguistic features, including n-gram, distributional semantic, and syntactic features, to detect the abusive language in user-generated online content. They also released a first of its kind corpus of user comments annotated for offensive language sampled from different domains.

Zampieri et al. (2019) proposed the Shared Task OffenseEval 2019 for detecting offensive language in the tweets sampled from Twitter. They released the *Offensive Language Identification Dataset (OLID)* consisting of 14,100 tweets in English annotated for offensive content using a fine-grained three-layer annotation scheme. OffenseEval 2020 Zampieri et al. (2020) uses the same dataset for the English task and similar data creation methodology for other languages.

Artetxe and Schwenk (2018) introduced an architecture for learning joint multilingual sentence representations, using a single BiLSTM encoder with a BPE vocabulary shared by all the languages. They use 93 languages, spanning across 30 different language families and 28 different scripts for training the sentence representation embeddings. Experimental results on the XNLI dataset (Conneau et al., 2018) for cross-lingual natural language inference and MLDoc dataset for cross-lingual document classification show that the cross-lingual transfer of learned features helps in improving the performance of the classification models in a cross-lingual setting.

Chen and Cardie (2018) proposed a fully unsupervised framework for learning multilingual word embeddings using only monolingual corpora. Unlike prior work in multilingual word embeddings, these embeddings exploit the relations between all the language pairs. By mapping all monolingual embeddings into a shared multilingual embedding space via a two-stage algorithm consisting of *Multilingual Adversarial Training* and *Multilingual Pseudo-Supervised Refinement*, the authors propose an effective method for learning the representation.

Recent works for the task of text classification exploit pretrained contextualized word representations rather than context-independent word representations. These pretrained contextualized word embeddings, like BERT (Devlin et al., 2018) and RoBERTa (Liu et al., 2020), have outperformed many existing techniques on most NLP tasks with minimal task-specific architectural changes and training resources. Contextualized word embeddings have been used for detecting subjective bias (Pant et al., 2020), and detecting hate speech in tweets (Mozafari et al., 2019).

To train multilingual variants of the contextualized word embeddings effectively, Lample and Conneau (2019) proposed a cross-lingual language modeling methodology termed as *XLM*. In their paper, they proposed two variants of their model: supervised and unsupervised multilingual language models. The supervised variant used parallel data with a new cross-lingual language model objective. Conneau et al. (2019) proposed *XLM-R* extending *XLM*. *XLM-R* differs from *XLM* in its ability to allow pretraining of the cross-lingual language models at scale using larger datasets and better-tuned hyperparameters. Their experiments on XNLI, MLQA, and NER show significant increases over the previous state-of-the-art. We further finetune *XLM-R*, pretrained on 2.5 TB Common Crawl corpus spanning 100 languages.

3 Methodology

In this section, we describe the preliminaries for the work: *XLM-R*, a cross-lingual contextual word representation, and cross-lingual inductive transfer. We then illustrate our sequential approach for finetuning the models.

3.1 XLM-R

XLM-R is a transformer-based cross-lingual model pretrained using a multilingual masked language model objective on 2.5 TB of CommonCrawl data in 100 languages. *XLM-R* obtains the state-of-the-

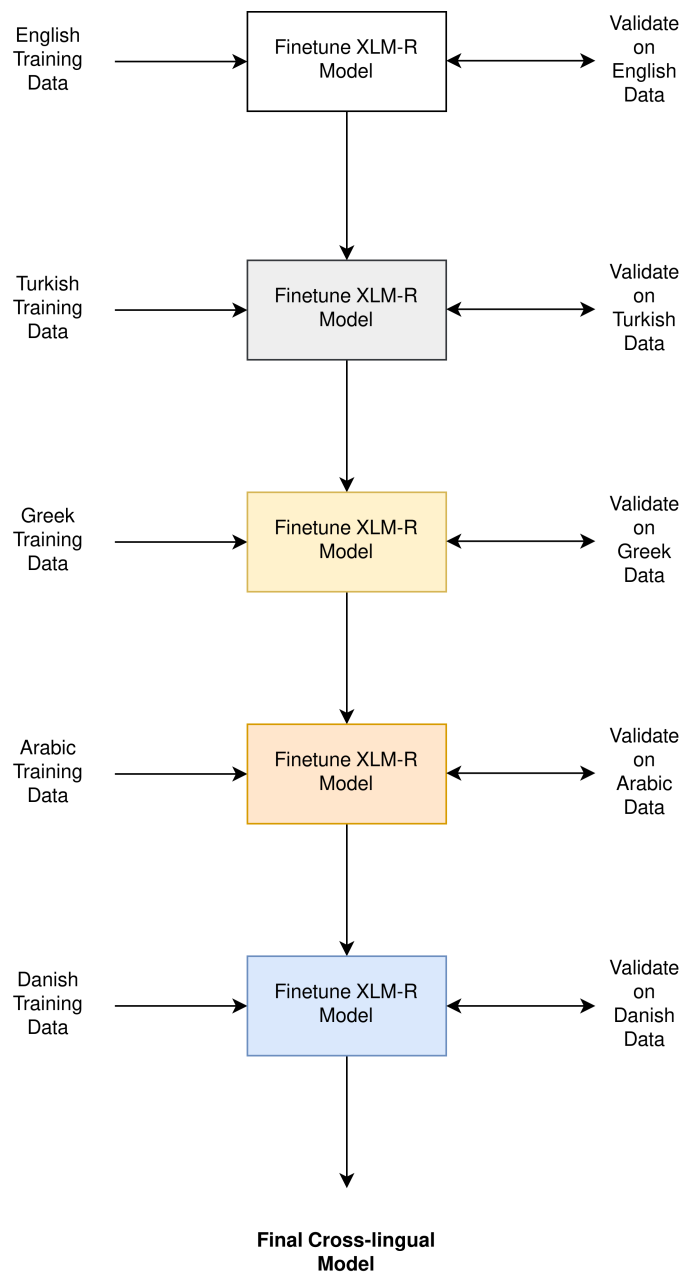


Figure 1: Architecture of our cross-lingual model.

art performance in cross-lingual classification, sequence labeling, and question answering. It obtains competitive results when compared with monolingual models on GLUE (Wang et al., 2018) and XNLI tasks, showing that it is possible to have a single large model for all languages without sacrificing per-language performance.

XLM-R, like other multilingual contextual word embeddings, suffers from the curse of multilinguality (Conneau et al., 2019). While low-resource language performances can be improved by adding high resource languages using pre-training, the overall downstream task may suffer from capacity dilution (Arivazhagan et al., 2019). This degradation happens because model capacity is constrained due to practical considerations including memory and speed during training. Moreover, for a fixed size model, the per-language capacity decreases with an increase in the number of languages. Despite these limitations, *XLM-R* provides competitive results on downstream tasks when compared with mono-lingual models.

3.2 Cross-lingual inductive transfer

Multilingual contextual word embeddings retain a partial level of alignment, as has been observed in (Cao et al., 2020). In *XLM* and *XLM-based* embeddings, multiple languages are specifically trained together with a sentence encoder leading to a higher degree of similarity in the alignment of corresponding words in different languages. Consequently, *XLM* word embeddings have shown to perform competitively in SemEval’17 cross-lingual word similarity task (Lample and Conneau, 2019). Thus, *XLM* achieves a higher degree of alignment while learning cross-lingual representation.

OffensEval 2020 provides a multilingual dataset for offensive language detection in five different languages, allowing one language’s learning to aid another. Therefore, cross-lingual inductive transfer learning using *XLM-R* with relatively aligned embeddings allows for the inductive transfer of linguistic features among different languages.

Figure 1 illustrates our chain-like model for detecting offensive detection on all five languages trained in a sequential manner. While fine-tuning for every language, we validate the model only with the current language upon which it is being fine-tuned.

4 Experiments and Results

In this section, we present details on the multilingual OffenseEval dataset, experimental settings for the validation experiments, and specifications of the system runs.

4.1 Dataset

In this subsection, we provide a comprehensive statistical analysis of the *multilingual Offensive Language Identification Dataset* (mOLID). It comprises tweets in five different languages: English (Rosenthal et al., 2020), Turkish (Çöltekin, 2020), Greek (Pitenis et al., 2020), Arabic (Mubarak et al., 2020), and Danish (Sigurbergsson and Derczynski, 2020). The tweets in English are annotated using a fine-grained three-layer annotation scheme, whereas the other four languages are annotated using a coarse-grained annotation scheme. For the English language, they use a three-level hierarchical annotation schema to distinguish between whether the language is offensive or not, the type of offensive language, and the target.

Statistical analysis of the dataset in Table 1 shows significant variation in the number of instances per language. Turkish has the maximum examples with 31,277 examples, whereas Danish has 2,960 examples, comprising 48.70% and 4.60% of total data, respectively. Further analysis shows that positive instances, signifying offensive languages, are significantly lesser than non-offensive sentences. Offensive sentences comprise only 23.21% of total data. This percentage, however, varies from language to language, ranging from 12.97% for Danish to 33.23% for English.

Language/Label	Positive	Negative	Total
English	4400	8840	13240
Turkish	6046	25231	31277
Greek	2486	6257	8743
Arabic	1589	6411	8000
Danish	384	2576	2960
Total	14905	49315	64220

Table 1: Number of instances with associated label and language.

4.2 Experimental Setting

In this subsection, we outline the experimental setup for the task and present the results obtained on both the validation dataset and the blind test set. For experiments, we used *XLM-R* having 16,550M parameters and 250K vocabulary size. For validation, we train and evaluate our model using 90 – 10

train-validation data split for the five languages in the given order: English, Turkish, Greek, Arabic, and Danish.

For models of all the languages, we finetune *XLM-R* with a learning rate of $1 * 10^{-5}$ for 2 epochs each with a maximum sequence length of 50 and a batch size of 32. We evaluate all the models on the following metrics for the binary classification: *F1*, *Precision*, *Recall*, *Accuracy*.

Language/Metric	F1 Score	Precision	Recall	Accuracy
English	0.709	0.728	0.690	0.808
Turkish	0.689	0.791	0.611	0.889
Greek	0.728	0.839	0.642	0.867
Arabic	0.729	0.762	0.698	0.907
Danish	0.676	0.686	0.667	0.915

Table 2: Experimental Results for Validation dataset for our final model.

Table 2 shows the performance of our proposed model on the validation set for the five languages during the fine-tuning of the cross-lingual contextual embeddings. We observe that the performance metrics obtained vary significantly with the languages, with *F1* varying from 0.676 in Danish to 0.729 in Arabic, and *Accuracy* varying from 0.808 in English to 0.915 in Danish.

Language/Metric	F1
English	0.919
Turkish	0.781
Greek	0.803
Arabic	0.848
Danish	0.759

Table 3: Experimental Results for Test dataset for our final model.

Table 3 shows the performance of our proposed model on the test dataset held out by the organisers for all the five languages. Organisers used *Macro-F1 score* for evaluating the model performance on test data. We observe that *F1-score* varies from language to language, from 0.919 from English to 0.759 for Danish.

5 Discussion

XLM-R is an unsupervised cross-lingual representation pretrained using transformers on an enormous scale across 100 languages. It can be fine-tuned for different downstream tasks in multiple languages, allowing one to take advantage of cross-lingual transfer learning. Like other transformer-based contextual word embeddings, *XLM-R* is highly scalable, and the process of fine-tuning takes minimum training efforts and resources.

Our proposed model, extending *XLM-R*, was trained in five different languages for the detection of offensive language. Owing to the training of *XLM-R* on hundred languages, our model allows a direct extension to other commonly-used languages. Our chain-like model exploits learnings from one language for detecting offensive language on another unrelated language.

We further performed experiments for detecting offensive language in a language different from the ones that we have trained on without supervision. This form of learning, known as zero-shot learning, has been demonstrated in recent related works (Artetxe and Schwenk, 2018; Lample and Conneau, 2019). To evaluate the efficacy of our model to perform zero-shot learning, we detect offensive language text in German using our final cross-lingual model. We perform the experiments in the same experimental setting as the GermEval Shared Task on the Identification of Offensive Language’s Subtask 1 (Wiegand et al., 2018), which entails coarse-grained binary classification of offensive language.

Model/Metric	F1	Precision	Recall	Accuracy
Paraschiv and Cercel (2019)	0.770	0.764	0.776	0.794
Our Model (Zero-Shot Learning)	0.700	0.795	0.682	0.781

Table 4: Experiments performed for German Offensive Language Detection demonstrating zero-shot learning.

Our experiments, illustrated in Table 4 show results competitive to the best-performing model by Paraschiv and Cercel (2019), which used supervised training data. Thus our model, despite not using supervised data in German, achieves competitive performance with a difference of 1.3% in *Accuracy*, further showing an increase in the metric of *Precision* by a margin of 4.06%.

6 Conclusion

This work presents a cross-lingual inductive transfer learning approach to detect the offensive language in tweets across five languages: English, Turkish, Greek, Arabic, and Danish. Our proposed model performs competitively for all five languages. We further performed comprehensive experiments to show that our model performs well in a zero-shot setting, which can be useful in the case of low resource languages. Our proposed model can be easily extended to other languages with minimum training efforts and resources, performing competitively even with significantly imbalanced data. Future work may involve including multiple unrelated languages to enable a universal offensive language detection model and the application of the proposed approach in other cross-lingual tasks.

References

- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, M. Gatu Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges. *ArXiv*, abs/1907.05019.
- Mikel Artetxe and Holger Schwenk. 2018. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. Multilingual alignment of contextual word representations. In *International Conference on Learning Representations*.
- Çağrı Çöltekin. 2020. A Corpus of Turkish Offensive Language on Social Media. In *Proceedings of the 12th International Conference on Language Resources and Evaluation*. ELRA.
- Xilun Chen and Claire Cardie. 2018. Unsupervised multilingual word embeddings. In *EMNLP*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, F. Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *ArXiv*, abs/1911.02116.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. In *NeurIPS*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Roberta: A robustly optimized bert pretraining approach. In *International Conference on Learning Representations*. accepted.
- Marzieh Mozafari, Reza Farahbakhsh, and Noel Crespi. 2019. A bert-based transfer learning approach for hate speech detection in online social media. In *COMPLEX NETWORKS*.

- 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*.
- Chikashi Nobata, Joel R. Tetreault, Achint Oommen Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *WWW*.
- Kartikey Pant, Tanvi Dadu, and Radhika Mamidi. 2020. Towards detection of subjective bias using contextualized word embeddings. In *Companion Proceedings of the Web Conference 2020*, WWW '20, page 75–76, New York, NY, USA. Association for Computing Machinery.
- Andrei Paraschiv and Dumitru-Clementin Cercel. 2019. Upb at germeval-2019 task 2: Bert-based offensive language classification of german tweets. 10.
- Zeses Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. Offensive Language Identification in Greek. In *Proceedings of the 12th Language Resources and Evaluation Conference*. ELRA.
- Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2020. A Large-Scale Semi-Supervised Dataset for Offensive Language Identification. In *arxiv*.
- Gudbjartur Ingi Sigurbergsson and Leon Derczynski. 2020. Offensive Language and Hate Speech Detection for Danish. In *Proceedings of the 12th Language Resources and Evaluation Conference*. ELRA.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium, November. Association for Computational Linguistics.
- 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*.
- Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012. Learning from bullying traces in social media. In *Proceedings of the 2012 conference of the North American chapter of the association for computational linguistics: Human language technologies*, pages 656–666. Association for Computational Linguistics.
- 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 *NAACL-HLT*.
- 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*.