



HAL
open science

Hate speech and offensive language detection using an emotion-aware shared encoder

Khouloud Mnassri, Praboda Rajapaksha, Reza Farahbakhsh, Noel Crespi

► **To cite this version:**

Khouloud Mnassri, Praboda Rajapaksha, Reza Farahbakhsh, Noel Crespi. Hate speech and offensive language detection using an emotion-aware shared encoder. IEEE International Conference on Communications (ICC), May 2023, Roma, Italy, France. 10.1109/ICC45041.2023.10279690. hal-04003849

HAL Id: hal-04003849

<https://hal.science/hal-04003849v1>

Submitted on 24 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hate Speech and Offensive Language Detection using an Emotion-aware Shared Encoder

Khouloud Mnassri, Praboda Rajapaksha, Reza Farahbakhsh, Noel Crespi

Samovar, Telecom SudParis, Institut Polytechnique de Paris, 91120 Palaiseau, France

{khouloud.mnassri, praboda_rajapaksha, reza.farahbakhsh, noel.crespi}@telecom-sudparis.eu

Abstract—The rise of emergence of social media platforms has fundamentally altered how people communicate, and among the results of these developments is an increase in online use of abusive content. Therefore, automatically detecting this content is essential for banning inappropriate information, and reducing toxicity and violence on social media platforms. The existing works on hate speech and offensive language detection produce promising results based on pre-trained transformer models, however, they considered only the analysis of abusive content features generated through annotated datasets. This paper addresses a multi-task joint learning approach which combines external emotional features extracted from another corpora in dealing with the imbalanced and scarcity of labeled datasets. Our analysis are using two well-known Transformer-based models, BERT and mBERT, where the later is used to address abusive content detection in multi-lingual scenarios. Our model jointly learns abusive content detection with emotional features by sharing representations through transformers’ shared encoder. This approach increases data efficiency, reduce overfitting via shared representations, and ensure fast learning by leveraging auxiliary information. Our findings demonstrate that emotional knowledge helps to more reliably identify hate speech and offensive language across datasets. Our hate speech detection Multi-task model exhibited 3% performance improvement over baseline models, but the performance of multi-task models were not significant for offensive language detection task. More interestingly, in both tasks, multi-task models exhibits less false positive errors compared to single task scenario.

Index Terms—Social media, Natural Language Processing, Hate speech, offensive language, Twitter, BERT, Multilingual BERT, Multi-task learning, emotional knowledge, shared encoder.

I. INTRODUCTION

People have become addicted to social media platforms in recent decades as means of engaging and connecting with each other. Through social media platforms like Twitter, individuals are increasingly communicating and expressing their opinions, and emotions. However, their content can contain harmful information prejudiced against a certain person or group, manifested as abusive language. It is challenging to have a concrete final definition of hate and offensive language, but in general, according to the United Nations¹: “hate speech” refers to **offensive** discourse targeting group or an individual based on inherent characteristics (race, religion or gender) and that may threaten social peace.” As for offensive language, it is more general referring to any content that can implicitly or explicitly offend the other or make him uncomfortable. Today, it becomes challenging and even

impossible to manually track the substance of posts due to the huge and unregulated content that is uploaded every day online. However, there have been many research attempts in automating the detection of abusive content on online platforms. The majority of these attempts adopt supervised learning and deep learning methods that were trained using an annotated dataset [1]. Among the substantial drawbacks in these approaches include the lack of training data, and in data bias, as well as the ambiguity of the abusive content that can be challenging to accurately detect using traditional NLP methods [2]. The latters start with using machine learning [2], then deep learning, transfer learning and ensemble learning [3]... Thus, we aim to involve training a single transformer-based model on multiple tasks simultaneously as it has been shown to improve performance of abusive language detection. This paper proposes a multi-task joint learning approach that utilizes additional features (emotions), to improve model performances. Emotion categorization from the text is deeply aligned with scientific concepts that have been studied for a very long time in the framework of Sentiment Analysis (SA) [4]. In fact, emotion classification seeks to automatically classify texts in a precise manner according to several classes such as anger, fear... [5]. A psychology study [6] found a direct correlation between the speaker’s psychological and emotional state, and his abusive speech [7]. For instance, abusive content presents unpleasant attitudes and feelings like anger, disgust, fear, and sadness. Hence, in this paper, we studied new methods of improving the detection of hate speech and offensive language by integrating emotional knowledge as an additional related feature. In the Multi-Task Learning (MTL) scenario, multiple tasks are learned in parallel while using a shared representation [8]. In comparison to learning multiple tasks individually and sequentially, this joint learning approach effectively increases the sample size while training a single model, which leads to improved performance by increasing its generalization [9]. Based on the recent state-of-the-art results given by implementing BERT in hate speech/offensive language detection [10]–[12], as well as mBERT (in cross-lingual setting) [4], [10], we have chosen to build our MTL model using those two language models. Following listed the main contributions of this paper:

- Building a multi-task learning framework that enables the model by sharing representations between several related tasks and generalizing better by achieving better

¹<https://www.un.org/en/hate-speech>

performance for the hate speech and offensive language detection task.

- Use of BERT and mBERT pre-trained models as shared encoders to create an MTL model, and analyze and compare their efficiency.
- Use of external related features (emotions) to improve hate speech detection task. We proved the hypothesis of the relationship between the spread of hateful and abusive content and the emotional psychological status of its writer.
- Model optimization within learning multiple tasks in parallel while using a shared transformer representation, which helps in avoiding the computation expenses and task-specific fine-tuning training step. This ensures making predictions at inference time faster than training two different models for every task independently.

Our joint learning approach uses a transformer-based shared encoder to implement a multi-task model for categorizing hate speech and offensive language. It aims to solve the issue of labeled data scarcity by sharing representations between several tasks, using an auxiliary dataset from the secondary task (emotion knowledge). The proposed multi-task model exhibits higher performance with fewer classification errors compared to single-task baseline models in both hate speech and offensive language detection.

II. LITERATURE SURVEY

A. Multi-task learning on hate/offensive speech

The implementation of multitask learning approach in the field of NLP abusive language detection remains a new approach. The first approaches used deep neural networks. Liu et al. [13] proposed a three-level framework. It includes detecting: hate speech, its types, and its topics. They developed a fuzzy ensemble approach in the setting of multi-task learning considering each type of hate speech as task prediction head. Their ensemble gave a detection rate of 0.93. Moreover, Kapil et al. [14] proposed a deep learning Shared-Private multi-task model to leverage the information from 5 abusive tasks. As a result, they built 4 deep neural networks, and training them on 5 datasets, they managed to get 26 models' combinations. Their approaches outperformed the single-task's with macro-F1 scores between 10% and 27%. Furthermore, Abu Farah et al. [15] worked on Arabic language within a joint learning approach, using sentiment analysis. Their best model is a multitask learning framework, based on CNN-BiLSTM. It gave macro F1-score of 0.9 and 0.7 for offensive and hate speech tasks respectively.

B. Multi-task learning based on Transfer learning:

With the outstanding performances given by the transformers, most of the researchers have employed these pre-trained models in hate speech detection. Awal et al. [16] proposed "AngryBERT", a BERT-based multitask model that learns with sentiment analysis and target detection as auxiliary tasks. They demonstrated the ability of those latters to improve hate speech detection. Their model gave

an F1 macro score of 90.71% on the Davidson dataset [2]. In addition, to address Aggression Identification, Samghabadi et al. [17] provided a neural model that builds attention on top of BERT using a multi-task learning paradigm. Their model scored 0.8579 weighted F1 on the English "Misogynistic Aggression Identification" task. Moreover, using the pre-trained AraBERT, Djandji et al. [18] enhanced this transformer with the inclusion of Multi-task learning to build a model able to learn well from little amounts of data, they trained it on several Arabic abusive speech datasets. Their model gave an F1 macro score of 90.15% and 83.41% in Offensive language and hate speech tasks respectively.

C. Emotion knowledge to detect hate/offensive language

In many fields, like the identification of mental diseases and social media analytics, understanding human emotional patterns is crucial. Since hate/offensive speech detection is integrally tied to the speaker's emotional state [6], detecting emotions has also become an essential application. Markov et al. [19] studied how stylometric and emotional characteristics affect the detection of hate speech. They demonstrated that, when integrated into an ensemble with deep learning models, the use of those features surpasses the commonly used ones to detect hateful content. Moreover, Chiril et al. [20] investigated the affective knowledge extracted from Sentic computing resources and from structured hate lexicons. They implemented multitasking techniques, and they attained the greatest outcomes with models that used data from these affective resources. Adding to that, Plaza-del-Arco et al. [21] defended the hypothesis of the relationship between hate speech tasks and sentiments, emotions as well as targets, via a straightforward multi-task learning architecture. Implementing BERT-based multitask model, they got the best overall result of $F1 = 0.79$. Furthermore, using a Transformer-based model, they also proposed a Multi-task model that makes use of shared sentiment and emotional knowledge to identify hate speech in Spanish tweets [22]. Their findings demonstrate that these knowledge work together to more reliably identify hate speech. Based on the above-mentioned works, we propose several strategies for the same tasks by integrating the best of these approaches. Thus, we built a BERT-based and mBERT-based multi-task model using the knowledge extracted from emotion samples provided by social media platforms using a large-size and more diverse-resources emotional dataset, and, implementing a shared representation to enable knowledge transfer between tasks and to reduce model complexity.

III. METHODOLOGY

A. Dataset

We conducted our experiments on Davidson dataset (Twitter) [2]: related to hateful and abusive language detection, and GoEmotion corpora (Reddit) [5]: related to the emotion analysis. The corpora statistics are displayed in Table I.

1) **Hate/Offensive speech dataset:** For the training of our model for the hate/Offensive task, we used Davidson corpora [2]. This data was compiled using a lexicon of hate speech

content taken from Twitter and classified tweets into *Hate speech*, *Offensive*, and *Neither* classes having about 24k total number of labeled samples. In this paper, we implement hate/offensive speech binary classification and thus, we created 2 corpora by separating the hateful labeled samples from offensive ones. As a result, we got Davidson-HATE and Davidson-OFF for hate speech and offensive language labeled datasets, respectively.

2) **Emotion dataset:** For training the emotion task, we use GoEmotions corpora [5], developed by Demzky et al. in 2020. This corpus is considered the largest manually annotated dataset, which consists of 58k English Reddit comments categorized as either Neutral or one of 27 emotion groups. The Ekman level further defines the emotion categorization into anger, disgust, fear, joy, sadness, and surprise. We implement the final corpora labeled as the Ekman model in our experiments in order to use a more generalizable dataset with less noisy data.

Dataset	Size	Labels
Davidson-HATE	5593	hate normal
Davidson-OFF	23 353	offensive normal
GoEmotions	48 834	anger, disgust, fear, joy, surprise, sadness, neutral

TABLE I

GENERAL OVERVIEW OF THE DATASETS ALONG WITH THEIR NUMBER OF SAMPLES AND LABELS’ DISTRIBUTIONS.

B. Transformer-based multi-task approach MTL

Unlike Single Task Learning (STL) which learns task-specific features from one dataset at a time, Multi-Task Learning (MTL) aims to tackle multiple issues at once. In STL, input vectors x_i are supervisedly mapped to any label y_i in order to train the model to complete a classification task T . Each sentence x_i is processed through the model layers, and the final representation is then run through Softmax to predict the probability distribution over C classes. Given a dataset D with n training labeled inputs (x_i, y_i) , the models’s weights are trained in order to reduce the cross-entropy of the detected \hat{y} and labeled y samples. Where,

$$\hat{y} = \text{softmax}(W + b) \quad (1)$$

Knowing that, W is the final weight after optimization of the linear classifier when training, and b is a bias term [14]. The objective of MTL scenario is to employ the method of learning numerous tasks to enhance performance on each one of them [8]. Although they could have different data or characteristics, these tasks are correlated and share some similarities. And when the model is trained, it can exploit shared characteristics by using some hints from one task to enhance the other ones. To understand more the process of building an MTL model, Zhang et al. [9] defined MTL as follows: Given n learning tasks, $\{T_i\}_{i=1}^n$, MTL seeks to enhance the learning of a model for the classification task T_i by leveraging the knowledge in some or all of the n tasks, where all or a subset of the tasks are connected. The two most popular methods to share knowledge in multi-tasking are hard parameter sharing and soft parameter

sharing [23]. **Hard parameter sharing** involves all tasks sharing the hidden layers with a number of task-specific output layers, which is the method we used by implementing shared BERT and mBERT encoders, in order to avoid task-specific parameters for each task, because it can lead to augmented model complexity [24]. On the other side, each task can have its own layers with certain shareable components, known as **soft parameters sharing**. Overall MTL approach efficiently improves the sample size when training a model, which leads to enhanced performance by improving the generalization of model in comparison to learning several tasks separately [9].

In this study, we used three related tasks: Hate speech detection, Offensive language detection, and Emotion recognition. The purpose is to determine whether implementing an MTL scenario to emotion categorization task facilitates the identification of hate/offensive speech, regardless of the source of social media data. Therefore, we build a typical contextualized embedding configuration where the input is represented by a well-known model, the Bidirectional Encoder Representations from Transformers (BERT) [25] and the Multilingual BERT (mBERT)². We used the latter model in order to build a cross-lingual generalizable approach, it could be tested on different target languages or used to be trained on monolingual low-resource ones. We used it also to compare with BERT based MTL model to understand more about the influence of adding features from other languages. We added two sequence classification heads to the encoder, one for Hate/offensive speech and another for Emotion recognition. The two tasks jointly share the transformer encoder, as seen in Figure 1, so that one task profits from the other by sharing features. The importance of using the shared encoder is to guarantee that any adjustments on its weights during training will change the same encoder weights, and not to use any extra GPU memory. The output heads for every task are then generated, and each task head is connected to a common sentence encoder. The layers are then adjusted in accordance with the specified set of our downstream tasks. As shown in Figure 1, the input representation is BERT/mBERT-based tokenization, and each task corresponds to a specific classification head. In the first step, given a data input, it is first tokenized using the default tokenizer of BERT/mBERT and then converted into pre-trained BERT embeddings: $Eb = \{eb_1, eb_2, \dots, eb_n\}$. These embeddings are then sent to the pre-trained BERT/mBERT shared encoder. After defining the feature extraction function using the corresponding tokenizer, we utilize the “*dataset.map*” method from the NLP package to apply this function to our data inputs. This NLP library effectively manages the mapping and caches the features. We constructed a “*MultitaskDataloader*” that combines many data loaders (built to load data of each task) into a single loader. The latter aims to sample, randomly, data from these data loaders, build a task batch and produce the associated task name (attached to each batch data). Overall,

²<https://huggingface.co/bert-base-multilingual-uncased>

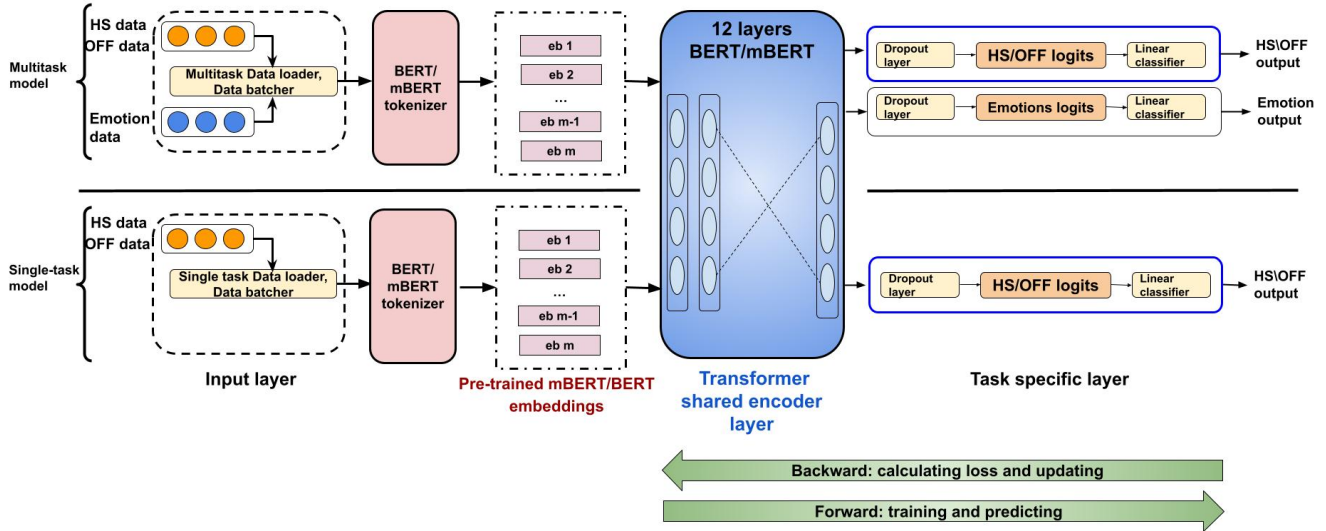


Fig. 1. Overview of the BERT and mBERT-based: Single-task, and Multitask leveraging emotions representations (auxiliary task) model architecture in the detection of Hate speech “HS” and Offensive language “OFF” (main tasks) from text input

information can flow from one task head to another through the shared encoder, which, is updated during the training via backpropagation. In fact, BERT/mBERT gets tuned by the combined loss of both tasks (Cross-Entropy loss) in order to learn a shared set of information between both tasks. As for the task-specific layers these consist of a task-specific softmax activation followed by a linear classification layer that is dedicated to extracting the unique information per task and giving final outputs.

IV. EXPERIMENTS AND RESULTS

A. Data preprocessing

We use the below procedures to pre-process the Twitter dataset using the Ekphrasis library [26]: 1) Switching into lowercase, 2) Delete URLs and emails, 3) Remove users’ names and mentions, 4) Shorten prolonged words and delete repeated characters (“yeeesss” to “yes”...), 5) Keep stop words, 6) Remove punctuations, unknown uni-codes, and additional delimiting characters, 7) Remove hashtags (#) and correct their texts (e.g, “#notracism” to “not racism”), 8) Eliminate tweets of length less than 2, and 9) Remove emojis.

B. Data analysis platform and evaluation metrics

The MTL models have been implemented using PyTorch. We trained the models on the training set and tested them on the validation set, keeping the same data split for GoEmotion corpora, and partitioning Davidson into 80% train and 20% validation set. The implemented models are trained using batch size 8 on Google Colab Pro (Tesla-T4 GPU environment with 32 GB of RAM). We used an optimizer with a learning rate of $1e-5$, and experimented with the Cross-entropy loss function. Since we used imbalanced datasets, classifiers’ performances are measured via multiple metrics: macro and weighted averaged F1 scores, precision, recall as well as accuracy. Weighted F1 calculates the score for each class and adds them together using a **weight** that depends on the number of true

labels of each class. Because we were using an imbalanced dataset, we want to assign more contributions to classes having more samples.

C. Results and interpretations

In this section, we compare the performance of single-task models with multi-task ones to understand the importance of using external, but related, features (emotions) in the abusive language detection. The STL models proposed in our previous contribution achieved considerable performances for detecting abusive content [12] where the ensemble average voting of BERT-CNN+BERT-LSTM gave better results in HS task and BERT-MLP (BERT-Multi Layer Perceptron) for the OFF task. Hence, we use these two best-performed models in [12], as well as BERT and mBERT as our single-task baseline models to compare with MTL approaches. Furthermore, we carried out an error analysis to get more information regarding the performance of the suggested MTL models. Working on imbalanced datasets, we want to get deeper into the classification of each class. So, we analyzed the confusion matrix and compared the MTL misclassification error with the other models as illustrated in the table II. Even though we have an emotion classifier as shown in Figure 1, this work mainly focuses on hate speech and offensive language detection tasks. Hence, Table II illustrates our experimental results for hate speech detection task “*HS task*” and offensive language detection task “*OFF task*”. In fact, the emotional analysis task trains the MTL network how to identify the emotion labels from the input samples and the representations generated by the encoder including the affective knowledge. This enables the MTL model to detect HS and OFF more accurately by leveraging the affective nature of the text input. Overall, the STL and MTL models’ results in both tasks (hate speech detection and offensive language detection) reveal

Model		Hate speech detection HS					Offensive language detection OFF				
		Acc.	Pr.	Recall	F1(m)	F1(w)	Acc.	Pr.	Recall	F1(m)	F1(w)
Baselines	BERT-based ensemble [12]	0.9474	0.9346	0.9234	0.9288	0.9470	-	-	-	-	-
	BERT-MLP [12]	-	-	-	-	-	0.9682	0.9400	0.9532	0.9465	0.9685
	BERT STL	0.9165	0.8956	0.8788	0.8867	0.9156	0.9689	0.9400	0.9566	0.9480	0.9692
	mBERT STL	0.9275	0.9109	0.8934	0.9016	0.9268	0.9645	0.931	0.952	0.9410	0.9649
BERT MTL		0.9385	0.9365	0.8971	0.9146	0.9371	0.9691	0.9367	0.9625	0.9489	0.9695
mBERT MTL		0.9413	0.9294	0.9123	0.9204	0.9407	0.9674	0.9346	0.9558	0.9459	0.9678

TABLE II

SINGLE-TASK LEARNING - STL VS MULTI-TASK LEARNING - MTL MODELS FOR HATE SPEECH (HS) AND OFFENSIVE LANGUAGE (OFF) DETECTION. BERT-based ensemble: Soft average voting Ensemble+BERT-CNN+BERT-LSTM, and MLP - Multi Layer Perceptron. Acc. - Accuracy and Pr. - Precision, F1(m) - F1 Macro and F1(w) - F1 Weighted.

a good performance when fine-tuning on small, imbalanced datasets.

1) *Hate speech detection*: As illustrated in Table II, for HS detection task, the performance of the mBERT-MTL model succeeded in surpassing all the STL models except the BERT-based ensemble model. Compared with the BERT STL model, these two models increased their performances by 3% and interestingly, there is no significant difference in their performances. For instance, the accuracy of the HS task for the ensemble model and mBERT-MLT model are 0.9474 and 0.9413, respectively and that of F1 macro are 0.9288 and 0.9204, respectively. Therefore, in order to understand the best model out of the BERT-based ensemble model vs mBERT MTL model in HS task, we conducted an error analysis based on the confusion matrix to explore which model exhibits less misclassification errors. Thus, we compare the confusion matrix of mBERT-MTL and BERT-based ensemble model as displayed in Figure 2. For HS task, although the ensemble model outperforms mBERT-MTL, we noted that the latter can detect Hate speech samples more efficiently than the Ensemble model, which is worth noticing since we used an imbalanced Hate speech dataset with only $\sim 77\%$ of samples labeled as “hate”. This indicates that the MTL model has less misclassification errors compared to the ensemble one, getting fewer percentages of false positives and false negatives for both classes.

2) *Offensive language detection*: We can interpret from Table II that, in OFF task, BERT-MTL gave the highest performances. However, we observe that, compared to the other mBERT-MTL model and the single-task models, the performance improvement is not very significant. The main reason behind this is due to the high imbalance Davidson-OFF corpora ($\sim 91\%$ offensive samples) [27]. Nevertheless, emotional features help to improve the OFF task classification by a small margin indicating that even with the highly imbalanced dataset, external features increase model performance. Further, we observed misclassification errors using the confusion matrix of BERT-MTL and BERT-STL models, the two best-performed models in the OFF task, as shown in Figure 3. BERT-MTL illustrates its ability to

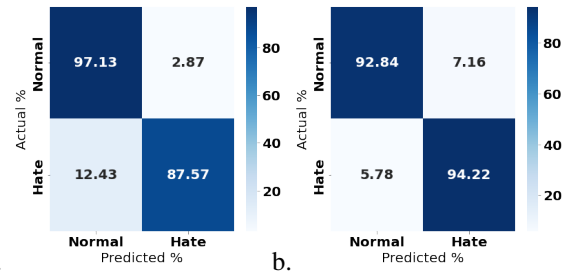


Fig. 2. Confusion matrix of Hate speech detection task: (a.) BERT-based ensemble model. (b.) mBERT multi-task model.

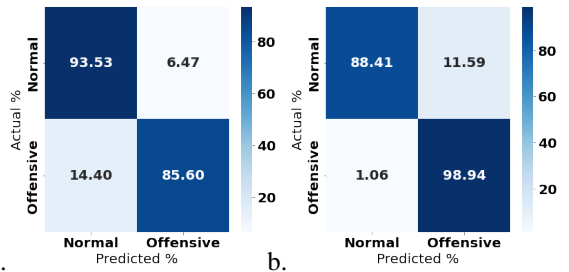


Fig. 3. Confusion matrix of Offensive language detection task: (a.) BERT single-task model, (b.) BERT multi-task model.

correctly detect offensive content with the largest rate among all baseline models, leading to True positives of 98.94% for “offensive” class, as well as the least misclassification error rate. In addition, the false positive rate of BERT-MTL model (1.06%) is fewer compared to BERT-STL model(14.40%). This indicates that multi-task models contributed to correctly classifying the texts in comparison to single-task models. Similar to the HS task, emotional features help to improve OFF task performance. Overall, based on the results obtained by the proposed MTL models for both hate speech and offensive language detection tasks, multi-task joint learning models outperformed single-task models and exhibit less false positive errors, indicating that emotional knowledge helps to improve the classification.

V. CONCLUSION AND FUTURE WORK

Recent years have seen a rise in the dissemination of abusive language, making it a significant issue for state governments

and social media corporations to find and delete this kind of content. As a result, we focus our paper on Hate speech and offensive language detection through a BERT/mBERT-based multi-task model to benefit from emotion analysis as a related task. Due to the sensitivity and granularity of hate speech and offensive language, we conducted experiments on two datasets extracted from Davidson corpora to consider, separately, hate speech and offensive language detection, each one, as a major classification task. The efficiency of our suggested model (in terms of performance and resource consumption) and a thorough investigation of the transfer of affective knowledge, demonstrate how emotion classification tasks enable the multi-task system to predict hate/offensive language more precisely by leveraging on this associated information. To improve the multi-task approach, we can focus more on the training datasets used, mainly by reducing the imbalance ratio through different data augmentation techniques to over-sample the corpora. As for the related external features, hate/offensive language detection task is not restricted only to emotion analysis, but can also be integrated with polarity, target (towards individuals or groups), irony, or sarcasm detection tasks. Thus we consider further using other related features. Furthermore, by implementing mBERT, and getting good performance, we can measure the cross-lingual generalization of our approach using zero-shot or few-shot learning and test it on several target low-resource languages (Arabic, French, German, etc.). We also aim to propose feature fusion approaches for hate speech detection as a joint learning approach using fuzzy ruling. In addition, we aim to compare different model performances in terms of resource consumption (i.e., memory, run-time) to determine the most optimized solution to deploy in real environments.

REFERENCES

- [1] Fortuna, T. & Nunes, S. 2018. A Survey on Automatic Detection of Hate Speech in Text. *ACM Comput. Surv.* **51**, 4, Article 85 (July 2019), 30 pages. <https://doi.org/10.1145/3232676>
- [2] Davidson, T., Warmusley, D., Macy, M. & Weber, I. Automated hate speech detection and the problem of offensive language. *Proceedings Of The International AAI Conference On Web And Social Media.* **11**, 512-515 (2017)
- [3] Jahan, M. S. & Oussalah, M. A systematic review of Hate Speech automatic detection using Natural Language Processing. *CoRR.* **abs/2106.00742** (2021), <https://arxiv.org/abs/2106.00742>
- [4] Medhat, W., Hassan, A. & Korashy, H. Sentiment analysis algorithms and applications: A survey. *Ain Shams Engineering Journal.* **5**, 1093-1113 (2014), <https://www.sciencedirect.com/science/article/pii/S2090447914000550>
- [5] Demszky, D., Movshovitz-Attias, D., Ko, J., Cowen, A., Nemade, G. & Ravi, S. GoEmotions: A Dataset of Fine-Grained Emotions. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics.* pp. 4040-4054 (2020,7)
- [6] Patrick, G. T. The psychology of profanity.. *Psychological Review.* **8**, 113 (1901)
- [7] Cao, R., Lee, R. K-W. & Hoang, T. DeepHate: Hate speech detection via multi-faceted text representations. *12th ACM Conference On Web Science.* pp. 11-20 (2020)
- [8] R. Caruana, "Multitask Learning," in *Learning to Learn*, S. Thrun and L. Pratt, Eds. Boston, MA: Springer US, 1998, pp. 95–133.
- [9] Zhang, Y. & Yang, Q. A Survey on Multi-Task Learning. *IEEE Transactions On Knowledge And Data Engineering.* **34**, 5586-5609 (2022)
- [10] Sohn, H. & Lee, H. MC-BERT4HATE: Hate Speech Detection using Multi-channel BERT for Different Languages and Translations. *2019 International Conference On Data Mining Workshops (ICDMW).* pp. 551-559 (2019)
- [11] Mozafari, M., Farahbakhsh, R. & Crespi, N. A BERT-Based Transfer Learning Approach for Hate Speech Detection in Online Social Media. *Complex Networks And Their Applications VIII.* pp. 928-940 (2020)
- [12] Mnassri, K., Rajapaksha, P., Farahbakhsh, R. & Crespi, N. BERT-based Ensemble Approaches for Hate Speech Detection. *IEEE GLOBECOM.* pp. 4649-4654 (2022)
- [13] Liu, H., Burnap, P., Alorainy, W. & Williams, M. L. Fuzzy Multi-Task Learning for Hate Speech Type Identification. *The World Wide Web Conference (WWW '19).* Association for Computing Machinery, New York, NY, USA, 3006–3012. pp. 3006-3012 (2019), <https://doi.org/10.1145/3308558.3313546>
- [14] Kapil, P. & Ekbal, A. A deep neural network based multi-task learning approach to hate speech detection. *Knowledge-Based Systems.* **210** pp. 106458 (2020), <https://www.sciencedirect.com/science/article/pii/S0950705120305876>
- [15] Abu Farha, I. & Magdy, W. Multitask Learning for Arabic Offensive Language and Hate-Speech Detection. *Proceedings Of The 4th Workshop On Open-Source Arabic Corpora And Processing Tools, With A Shared Task On Offensive Language Detection.* pp. 86-90 (2020,5), <https://aclanthology.org/2020.osact-1.14>
- [16] Awal, M. R., Cao, R., Lee, R. K-W. & Mitrović, S. AngryBERT: Joint Learning Target and Emotion for Hate Speech Detection. *Advances In Knowledge Discovery And Data Mining.* pp. 701-713 (2021)
- [17] Safi Samghabadi, N., Patwa, P., PYKL, S., Mukherjee, P., Das, A. & Solorio, T. Aggression and Misogyny Detection using BERT: A Multi-Task Approach. *Proceedings Of The Second Workshop On Trolling, Aggression And Cyberbullying.* pp. 126-131 (2020,5), <https://aclanthology.org/2020.trac-1.20>
- [18] Djandji, M., Baly, F., Antoun, W. & Hajj, H. Multi-Task Learning using AraBert for Offensive Language Detection. *Proceedings Of The 4th Workshop On Open-Source Arabic Corpora And Processing Tools, With A Shared Task On Offensive Language Detection.* pp. 97-101 (2020,5), <https://aclanthology.org/2020.osact-1.16>
- [19] Markov, I., Ljubešić, N., Fišer, D. & Daelemans, W. Exploring Stylometric and Emotion-Based Features for Multilingual Cross-Domain Hate Speech Detection. *Proceedings Of The Eleventh Workshop On Computational Approaches To Subjectivity, Sentiment And Social Media Analysis.* pp. 149-159 (2021,4)
- [20] Chiril, P., Pamungkas, E. W., Benamara, F., Moriceau, V. & Patti, V. Emotionally informed hate speech detection: a multi-target perspective. *Cognitive Computation.* pp. 1-31 (2022)
- [21] Arco, F. M. P., Halat, S., Padó, S. & Klinger, R. Multi-Task Learning with Sentiment, Emotion, and Target Detection to Recognize Hate Speech and Offensive Language. *CoRR.* **abs/2109.10255** (2021), <https://arxiv.org/abs/2109.10255>
- [22] Plaza-Del-Arco, F. M. D., Molina-González, M., Ureña-López, L. A. & Martín-Valdivia, M. T. A Multi-Task Learning Approach to Hate Speech Detection Leveraging Sentiment Analysis. *IEEE Access.* **9** pp. 112478-112489 (2021)
- [23] Xue, Y., Liao, X., Carin, L. & Krishnapuram, B. Multi-Task Learning for Classification with Dirichlet Process Priors. *Journal Of Machine Learning Research.* **8**, 35-63 (2007), <http://jmlr.org/papers/v8/xue07a.html>
- [24] Liu, Y., Zhuang, B., Shen, C., Chen, H. & Yin, W. Training Compact Neural Networks via Auxiliary Overparameterization. *CoRR.* **abs/1909.02214** (2019), <http://arxiv.org/abs/1909.02214>
- [25] Devlin, J., Chang, M., Lee, K. & Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR.* **abs/1810.04805** (2018), <http://arxiv.org/abs/1810.04805>
- [26] Baziotis, C., Pelekis, N. & Doukeridis, C. DataStories at SemEval-2017 Task 4: Deep LSTM with Attention for Message-level and Topic-based Sentiment Analysis. *Proceedings Of The 11th International Workshop On Semantic Evaluation (SemEval-2017).* pp. 747-754 (2017,8)
- [27] Chen, S., Zhang, Y. & Yang, Q. Multi-Task Learning in Natural Language Processing: An Overview. *CoRR.* **abs/2109.09138** (2021), <https://arxiv.org/abs/2109.09138>