

# Garain at SemEval-2020 Task 12: Sequence based Deep Learning for Categorizing Offensive Language in Social Media

**Avishek Garain**

Computer Science and Engineering

Jadavpur University

Kolkata, India

avishekgarain@gmail.com

## Abstract

SemEval-2020 Task 12 was OffenseEval: Multilingual Offensive Language Identification in Social Media (Zampieri et al., 2020). The task was subdivided into multiple languages and datasets were provided for each one. The task was further divided into three sub-tasks: offensive language identification, automatic categorization of offense types, and offense target identification. I have participated in the task-C, that is, offense target identification. For preparing the proposed system, I have made use of Deep Learning networks like LSTMs and frameworks like Keras which combine the bag of words model with automatically generated sequence based features and manually extracted features from the given dataset. My system on training on 25% of the whole dataset achieves macro averaged f1 score of 47.763%.

## 1 Introduction

The ever-growing amount of user-generated data on social media platforms be it Facebook, Twitter, blogs or any other electronic medium introduces new challenges in terms of automatic content moderation, especially regarding hate speech and offensive language detection. Not only is hate speech more likely to happen on the Internet, where anonymity is easily obtained and speakers are psychologically distant from their audience, but its online nature also gives it a far-reaching and determinative impact (Shaw, 2011). User content mostly consists of microposts, where the context of a post can be missing or inferred only from current events. Manual verification of each posting by a human moderator is infeasible due to the high amount of postings created every day. Consequently, automated detection of such attacking postings is the only feasible way to counter this kind of hostility. However, this task is challenging because natural language is fraught with ambiguities, and language in social media is extremely noisy. The classification system that would be prepared for the task, needed to be generalized for various test corpora as well. In this paper I have described the system consisting of a sequential pipeline with text feature extraction and classification as its main components. Firstly, a bag-of-words model is used for encoding the sentences into corresponding integer sequence. Thereafter, vectors are generated from these sequences and fed to a series of BiLSTM layers for training. Then a softmax layer is used for ternary classification into the corresponding offensive language categories.

The rest of the paper has been organized as follows. Section 2 describes the data, on which, the task was performed. The methodology followed is described in Section 3. This is followed by the results and concluding remarks in Section 4 and 5 respectively.

## 2 Data

The dataset that has been used to train and validate the model is an annotation following the tagset proposed in the Offensive Language Identification Dataset (OLID) (Zampieri et al., 2019a) and has been used in OffenseEval 2019. This year more accurate and quantified dataset was generated and annotated using a

---

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

<https://sites.google.com/site/offensevalsharedtask/offenseval2019>

<https://sites.google.com/site/offensevalsharedtask/home>

semisupervised approach to increase the quality and quantity of the dataset named the Semi-Supervised Offensive Language Identification Dataset (SOLID)(Rosenthal et al., 2020). The resulting dataset that is the newer version of the OLID dataset has been used for training and validating purposes. It was collected from Twitter; the data being retrieved the data using the Twitter API by searching for keywords and constructions that are often included in offensive messages. The vast majority of content is not offensive so different strategies have been tried to keep a reasonable number of annotations in the offensive class amounting to around 26% of the dataset.

The dataset provided consisted of annotations in their original form along with the corresponding labels. Subtask C consisted of the labels IND, GRP and OTH.

Label	Meaning
IND	Offensive tweet targeting an individual
GRP	Offensive tweet targeting a group
OTH	Offensive tweet targeting neither group or individual

Table 1: Meaning of the labels used in the dataset

The dataset has 188742 instances. Random selection of 50000 instances have been made which are divided into 35000 training data instances and 15000 validation data instances. The test data comprises of 850 instances.

Label	Train	Validation	Total
IND	28319	12136	40455
OTH	2062	884	2946
GRP	4619	1980	6599
All	35000	15000	50000

Table 2: Distribution of the labels in the dataset

Label	Test
IND	580
OTH	80
GRP	190
All	850

Table 3: Distribution of the labels in Test dataset

### 3 Methodology

My approach is to preprocess the annotations and then convert the annotations into a sequence of words and convert them into word embeddings. I then run a neural-network based algorithm on the processed tweet. Label based categorical division of data is given in Table 2. I have used SenticNet5(Cambria et al., 2018) for finding sentiment values of individual words. The sentiment features play a vital role in context of offensive language as it is during sad state of mind and hatred towards someone that offensiveness is at its peak. The Sentiment values range from -1 to 1 depicting Negative, Neutral and Positive sentiments. The use of BiLSTM networks in our model might have resulted in better results. The work done by Sepp Hochreiter et al. (Hochreiter and Schmidhuber, 1997) led the road-map to newer domain of work by bringing the concept of memory into usage for sequence based learning problems.

I first have taken the annotations and sent the raw data through some pre-processing steps, for which I took inspiration from the work on Hate Speech against immigrants in Twitter(Garain and Basu, 2019a), part of SemEval2019. The steps used here are built as an advancement of this work.

It consisted of the following steps:

1. Replacing emojis and emoticons by their corresponding meanings(Garain, 2019)
2. Removing mentions

3. Replacing words like "XX" or "XXX" with "sexual"
4. Contracting whitespace
5. Removing URLs

In step 1, for example,  
 ",-)" is replaced by "winking happy"  
 ":-C" is replaced with "real unhappy"  
 ";-(" is replaced with "crying"

Similarly I replaced 110 emoticons by their feelings.

The step 3 consists of manually identifying a feature of inclusion of words "XX" or "XXX" in sexual harassment cases and taking advantage of this identified feature. Using this extraction, contributed to increasing accuracy in classification to some extent. They played an important role during the model-training stage.

The pre-processed annotations are treated as a sequence of words with interdependence among various words contributing to its meaning. I then take a Bag-of-words model approach as well as use pre-trained GloVe vectors. For the Bag-of-words approach I convert the annotations into one-hot vectors as shown in Fig. 1. The Bag-of-words approach outperformed so I have shown results related to this encoding of text.

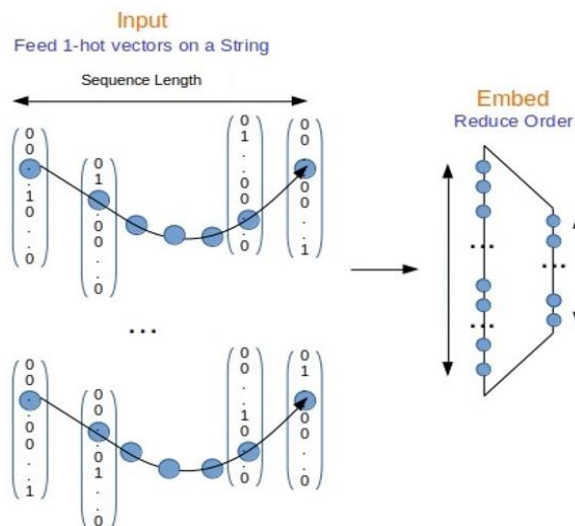


Figure 1: Bag-of-words model

Singular verbs clearly depict individuals while plural verbs depict groups and others, so their counts are included as features along with other manually extracted features listed below:

1. Counts of words with positive sentiment, negative sentiment and neutral sentiment in English
2. Frequency of auxiliary verbs like "is", "was", "are", "were"
3. Subjectivity score of the tweet (calculated using predefined libraries)
4. Frequency of difficult and easy words(Basu et al., 2019)
5. Number of question marks, exclamations and full-stops in the tweet
6. Frequency of words like "they", "he", "she", "we"

My model for the sub-task is a neural-network based model. For the task, first, I have merged the manually-extracted features with the feature vector obtained after converting the processed tweet to one-hot encoding. The output is processed through an embedding layer which transformed the tweet into a 128 length vector. The embedding layer maps the unique indices in the one-hot vector to the embedding vector space. I pass the embeddings through a Bidirectional LSTM layer containing 256 units. This is

followed by another bidirectional LSTM layer containing 512 units with its dropout and regular dropout set to 0.45 and activation being a sigmoid activation. This is followed by a Bidirectional LSTM layer with 128 units for better learning followed by a Dense Layer. The Dense layer consists of 3 units for three classes representing the classes Individual, Group and Others. The softmax layer gives a probability prediction percentage against each of the classes and thus finally gives the final results by assigning the label with maximum probability as output. The model is compiled using the Nadam optimization algorithm with a learning rate of 0.001. Categorical crossentropy is used as the loss function.

The architecture is depicted in Table 4.

Layer (type)	Output Shape	Param #
embedding_1 (Embedding)	(None, 100, 128)	6400000
bidirectional_1 (Bidirection)	(None, 100, 256)	263168
bidirectional_2 (Bidirection)	(None, 100, 512)	1050624
dropout_1 (Dropout)	(None, 100, 512)	0
bidirectional_3 (Bidirection)	(None, 256)	656384
dense_1 (Dense)	(None, 3)	514
Total params		8,370,690
Trainable params		8,370,690
Non-trainable params		0

Table 4: Neural Network Model Architecture

For the neural models in the language context, most popular are LSTMs (Long short term memory) which are a type of RNN (Recurrent neural network), which preserve the long term dependency of text. I use a Bidirectional-LSTM based approach to capture information from both the past and future context thus enhancing the context grasping capacity of the architecture. Figure 2 shows working of a BiLSTM unit.

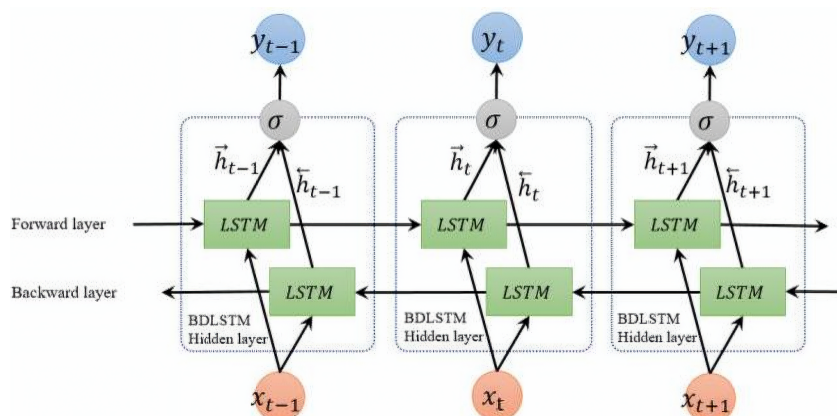


Figure 2: BiLSTM Unit

The dataset is skewed in nature. If trained on the entire dataset without any validation, the model tends to completely overfit to the class with higher frequency as it leads to a higher training accuracy score. To overcome this problem, I have taken some steps. Firstly, the training data has already been splitted into

two parts - one for training and one for validation thus, solving overfitting to some extent. The training is stopped when two consecutive epochs increased the measured loss function value and decrease in Validation accuracy for the validation set. Secondly, class weights have been assigned to the different classes present in the data which are chosen to be proportional to the inverse of the respective frequencies of the classes. Hypothetically, the model then gives equal weight to the skewed classes and this penalizes tendencies to overfit to the data.

#### 4 Results

I participated in subtask C of OffenseEval 2020 which is task 12 of SemEval 2020 and my system worked quite well. I have included the automatically generated macro averaged evaluation metrics along with the detailed metrics calculated using the gold labels. The results are depicted in Tables 5-6.

System	Train (%)	Validation (%)
Without	91.37	86.23
With	93.7	89.75

Table 5: Comparison of development phase accuracies with and without pre-processing operations

#### Task-C

Class	Precision	Recall	F1-score	support
GRP	0.69	0.44	0.54	190
IND	0.76	0.95	0.85	580
OTH	0.33	0.03	0.05	80
<b>macro avg</b>	<b>0.60</b>	<b>0.47</b>	<b>0.4776</b>	<b>850</b>

Table 6: Result Metrics calculated using Gold labels dataset

The confusion matrix is shown in Fig. 3.

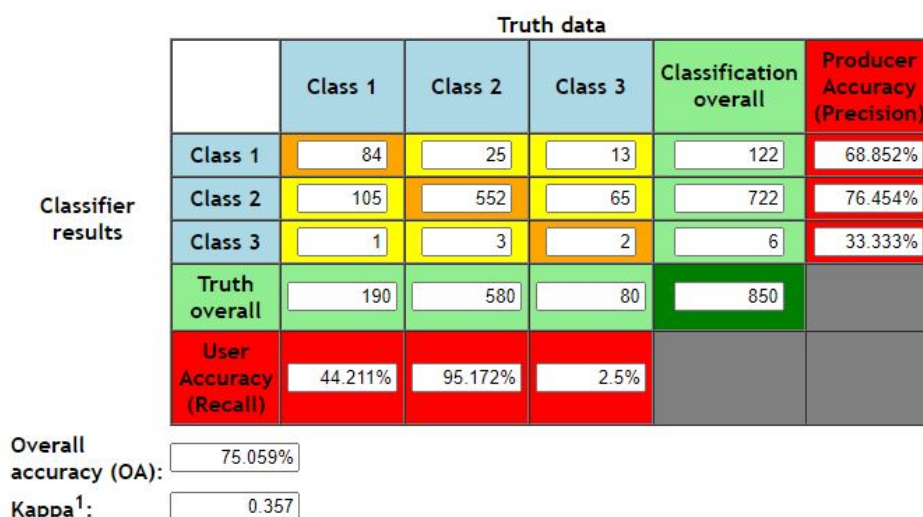


Figure 3: Confusion Matrix and Kappa value

The skewness in the data for the OTH class has led to systematic ignorance of the system to learning features related to the class thus resulting in worse results compared to other class labels of Offensiveness.

#### 5 Conclusion

In this paper, I have presented a model which performs satisfactorily in the given task. The model is based on a many2one sequence learning based architecture. There is scope for improvement by including more manually extracted features (like those removed in the pre-processing step) to increase the performance. I

could use only 25% of the whole dataset due to lack of computational resources. Data required for a Deep Learning model is quite high. Using the whole database would surely give excellent results. Removing the data constraints might lead to better accuracies and f1 metrics. Use of regularizers can also lead to proper generalization of model, henceforth increasing the metrics.

## References

- 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 the 13th International Workshop on Semantic Evaluation*, pages 54–63.
- A. Basu, A. Garain, and S. K. Naskar. 2019. Word difficulty prediction using convolutional neural networks. In *TENCON 2019 - 2019 IEEE Region 10 Conference (TENCON)*, pages 1109–1112.
- 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):223–242.
- Erik Cambria, Soujanya Poria, Devamanyu Hazarika, and Kenneth Kwok. 2018. Senticnet 5: Discovering conceptual primitives for sentiment analysis by means of context embeddings. In *AAAI*.
- Zhiyong Cui, Ruimin Ke, Ziyuan Pu, and Yin Hai Wang. 2018. Deep bidirectional and unidirectional lstm recurrent neural network for network-wide traffic speed prediction.
- Maral Dadvar, Dolf Trieschnigg, Roeland Ordelman, and Franciska de Jong. 2013. Improving cyberbullying detection with user context. In *Advances in Information Retrieval*, pages 693–696. Springer.
- Thomas Davidson, Dana Warmsley, Michael Macy, and Ingmar Weber. 2017. Automated Hate Speech Detection and the Problem of Offensive Language. In *Proceedings of ICWSM*.
- Karthik Dinakar, Roi Reichart, and Henry Lieberman. 2011. Modeling the detection of textual cyberbullying. In *The Social Mobile Web*, pages 11–17.
- Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan Bhamidipati. 2015. Hate speech detection with comment embeddings. In *Proceedings of the 24th International Conference on World Wide Web Companion*, pages 29–30. International World Wide Web Conferences Steering Committee.
- Mai ElSherief, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding. 2018. Hate Lingo: A Target-based Linguistic Analysis of Hate Speech in Social Media. *arXiv preprint arXiv:1804.04257*.
- Association for Computing Machinery. 1983. *Computing Reviews*, 24(11):503–512.
- Paula Fortuna and Sérgio Nunes. 2018. A Survey on Automatic Detection of Hate Speech in Text. *ACM Computing Surveys (CSUR)*, 51(4):85.
- Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large Scale Crowdsourcing and Characterization of Twitter Abusive Behavior. *arXiv preprint arXiv:1802.00393*.
- Björn Gambäck and Utpal Kumar Sikdar. 2017. Using Convolutional Neural Networks to Classify Hate-speech. In *Proceedings of the First Workshop on Abusive Language Online*, pages 85–90.
- Avishek Garain and Arpan Basu. 2019a. The titans at semeval-2019 task 5: Detection of hate speech against immigrants and women in twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 494–497.
- Avishek Garain and Arpan Basu. 2019b. The titans at semeval-2019 task 6: Offensive language identification, categorization and target identification. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 759–762.
- Avishek Garain and Sainik Kumar Mahata. 2019. Sentiment analysis at sepln (tass)-2019: Sentiment analysis at tweet level using deep learning.
- Avishek Garain, Arpan Basu, Rudrajit Dawn, and Sudip Kumar Naskar. 2019. Sentence simplification using syntactic parse trees. In *2019 4th International Conference on Information Systems and Computer Networks (ISCON)*, pages 672–676. IEEE.

- A. Garain, S. K. Mahata, and S. Dutta. 2020. Normalization of numeronyms using nlp techniques. In *2020 IEEE Calcutta Conference (CALCON)*, pages 7–9.
- Avishek Garain. 2019. Humor analysis based on human annotation (haha)-2019: Humor analysis at tweet level using deep learning.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November.
- Ritesh Kumar, Atul Kr. Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking Aggression Identification in Social Media. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC)*, Santa Fe, USA.
- Irene Kwok and Yuzhou Wang. 2013. Locate the hate: Detecting Tweets Against Blacks. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- 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 the 11th Forum for Information Retrieval Evaluation*, pages 14–17.
- Hamdy Mubarak, Darwish Kareem, and Magdy Walid. 2017. Abusive Language Detection on Arabic Social Media. In *Proceedings of the Workshop on Abusive Language Online (ALW)*, Vancouver, Canada.
- 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 Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive Language Detection in Online User Content. In *Proceedings of the 25th International Conference on World Wide Web*, pages 145–153. International World Wide Web Conferences Steering Committee.
- 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. *arXiv preprint arXiv:2004.14454*.
- Björn Ross, Michael Rist, Guillermo Carbonell, Benjamin Cabrera, Nils Kurowsky, and Michael Wojatzki. 2016. Measuring the Reliability of Hate Speech Annotations: The Case of the European Refugee Crisis. In *Proceedings of the Workshop on Natural Language Processing for Computer-Mediated Communication (NLP4CMC)*, Bochum, Germany.
- Anna Schmidt and Michael Wiegand. 2017. A Survey on Hate Speech Detection Using Natural Language Processing. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media. Association for Computational Linguistics*, pages 1–10, Valencia, Spain.
- LaShel Shaw. 2011. Hate speech in cyberspace: Bitterness without boundaries. *Notre Dame JL Ethics & Pub. Pol’y*, 25:279.
- 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.
- Zeeraq Waseem, Thomas Davidson, Dana Warmsley, and Ingmar Weber. 2017. Understanding Abuse: A Typology of Abusive Language Detection Subtasks. In *Proceedings of the First Workshop on Abusive Language Online*.
- 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. 2019a. Predicting the Type and Target of Offensive Posts in Social Media. In *Proceedings of NAACL*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019b. Predicting the Type and Target of Offensive Posts in Social Media. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 1415–1420.

Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019c. SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval). In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.

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*.

Ziqi Zhang, David Robinson, and Jonathan Tepper. 2018. Detecting Hate Speech on Twitter Using a Convolution-GRU Based Deep Neural Network. In *Lecture Notes in Computer Science*. Springer Verlag.