Stance Detection in Code-Mixed Hindi-English Social Media Data using Multi-Task Learning

Sushmitha Reddy Sane*[1] **Suraj Tripathi*[2]** **Koushik Reddy Sane[1]** **Radhika Mamidi[1]**
[1]International Institute of Information Technology, Hyderabad
[2]Indian Institute of Technology, Delhi
{sushmithareddy.sane, koushikreddy.sane}@research.iiit.ac.in,
surajtripathi93@gmail.com, radhika.mamidi@iiit.ac.in

Abstract

Social media sites like Facebook, Twitter, and other microblogging forums have emerged as a platform for people to express their opinions and views on different issues and events. It is often observed that people tend to take a stance; in favor, against or neutral towards a particular topic. The task of assessing the stance taken by the individual became significantly important with the emergence in the usage of online social platforms. Automatic stance detection system understands the user's stance by analyzing the standalone texts against a target entity. Due to the limited contextual information a single sentence provides, it is challenging to solve this task effectively. In this paper, we introduce a Multi-Task Learning (MTL) based deep neural network architecture for automatically detecting stance present in the code-mixed corpus. We apply our approach on Hindi-English code-mixed corpus against the target entity - "Demonetisation." Our best model achieved the result with a stance prediction accuracy of 63.2% which is a 4.5% overall accuracy improvement compared to the current supervised classification systems developed using the benchmark dataset for code-mixed data stance detection.

1 Introduction

The amount of data that is being generated by Internet users is massive and is multiplying every day. On the social media platform Twitter alone, users send more than 300k tweets per minute[2]. Users express their feelings, views and share their opinions on different topics ranging from politics, sports, government policies, movies, social issues, etc. More often, we observe that users tend to take a stance on a particular topic. Stance is a position on a specific issue, based on consideration of the evidence, often expressed publicly. It is an unpractical task to manually detect the stance represented by the individuals in these texts. The problem of automatic stance detection has caught the attention of researchers to effectively identify the stance taken by the user in numerous texts towards a particular topic.

1.1 Stance Detection

Stance detection addresses the problem of determining whether the author of a text is in FAVOUR of (positive), is AGAINST (negative) or is NEUTRAL (none) towards a particular target topic. The task of detecting stance closely compliments the task of sentiment analysis but is distinctive in nature (Mohammad, 2016). Stance detection considers the authors evaluative outlook towards specific targets rather than merely considering speakers emotions which adds to the problem of sentiment analysis.

1.2 Code-Mixing

The majority of the work in detecting stance has been done in English and other monolingual languages only. Our work focuses on code-mixed Hindi-English texts from users majorly in the Indian Subcontinent. It is improvisation to the task of detecting stance presented (Swami et al., 2018) for the target entity - i.e., Notebandi (Demonetisation), which was implemented in India. The government announced the issuance of new 500 and 2000 banknotes by exchanging with the demonetised notes. This action was taken to curb counterfeit cash used to fund terror groups. Many citizens of India and other nations, voiced their opinions and took a stance on this move by the Government of India.

Example: "*Demonetisation is a step towards the development and betterment of society*."

* These authors contributed equally to this work.
[2]http://www.internetlivestats.com/twitter-statistics/

Proceedings of the 10th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, pages 1–5
Minneapolis, June 6, 2019.

In this tweet, we can observe that the user most likely is in favor of the move. Our model for stance detection determines the stance taken by the tweeter automatically. An example of a tweet in the code-mixed Hindi-English corpus is

Example: "*Notebandi ne foreigners ko bhi pareshan karke rakha hai Demonetisation* ."

Here, the words *demonetisation*, *foreigners* are English while the others are Hindi. This sentence is transliterated into Hindi and then translated to English for employing English-based word representations.

In this paper, we describe an MTL based framework which makes use of deep learning architecture for automatic stance detection on social media corpus presented by (Swami et al., 2018). One of the major limitations in social media corpus is that users use unstructured text formats, non-grammatical structures and express rather explicitly compared to opinion surveys or formal texts. These informal usages introduce noise in the corpus and make the task very challenging. Also, the code-mixed corpus lacks the presence of word embeddings, commonly used, to train any deep learning model. So, we use machine transliterated and translated English corpus to feed to the network in order to use word2vec (Mikolov et al., 2013) based word embeddings.

The paper is organized as follows. In Section 2, we review related research in the area of stance detection and code mixing. In Section 3, we describe our system architecture to detect stance. In Section 4, we present the results and discuss the evaluation metrics. Finally, we conclude our work in Section 5 followed by future work in Section 6.

2 Related Work

Stance Detection problem is widely discussed and studied for the past few years in opinion mining. One of the initial work on stance classification (Somasundaran and Wiebe, 2010) explores the use of sentiment and arguing features for classifying stances in ideological debates by constructing an arguing lexicon from a manually annotated corpus. The combination of opinion target pair features was employed for the classification task. Later, Anand et al. (2011) identifies that for a particular topic, classification results using lexical and contextual features are far better than the best feature set without any contextual features analyzing the dialogic structure of debates. Walker et al. (2012); Hasan and Ng (2013) studied stance detection in two-side online debate data, and Faulkner (2014) examined document-level argument stance in student essays where the language of the texts are structured, monolingual and grammatically correct. And lately, a shared task for stance detection research focused on Twitter data (Mohammad et al., 2016).

Stance at user-level (Rajadesingan and Liu, 2014) is determined based on the assumption that if several users retweet one pair of tweets about a controversial topic, it is likely that they support the same side of a debate. Djemili et al. (2014) uses a set of rules based on the syntax and discourse structure of the tweet to identify tweets that contain ideological stance. However, none of these works attempts to determine the stance from a single tweet. In the field of social media mining, Guellil and Boukhalfa (2015) described in detail about different works in opinion mining and sentiment analysis and identified a set of open issues. Apart from English language, stance detection is carried out on Czech news commentaries (Krejzl et al., 2017) where maximum entropy classifier approach was used which were initially developed to detect stance in English tweets which uses sentiment and domain-specific features. Also, for the corpus of Spanish tweets (Anta et al., 2013), topic detection, and sentiment analysis approaches are used.

Multi-task learning approach (MTL) jointly trains multiple tasks in parallel, which acts as additional regularization, to improve the underlying network's generalization across all the tasks. It has proven to be a novel and effective learning schema in many NLP problems. Recently, multi-task learning approaches have been used for sentiment and sarcasm detection in (Majumder et al., 2019) , implicit discourse relationship identification (Lan et al., 2017), key-phrase boundary classification (Augenstein and Søgaard, 2017), improving sequence tagging tasks (Changpinyo et al., 2018) and improving named entity recognition tasks (Pham et al., 2019) and target dependent sentiment analysis (Gupta et al., 2019).

3 Method Description

The following subsections explain the preprocessing of the corpus and the deep learning architecture proposed for stance detection.

3.1 Preprocessing

Preprocessing is done on the tweets by removing twitter handles starting with "@" or words that had any special symbol. The word "Notebandi" is replaced by the phrase "noton par prathibandh." Emoticons have been removed, and URLs are replaced with the word "URL." This cleaned corpus is transliterated and translated into English sentences using Google translate API which is later given as input to the model.

3.2 Model Architecture

We propose a multi-channel convolutional neural network (CNN), refer Figure 1, for detecting stance from the given input text. Mutli-channel CNNs are used to expand the network in width without increasing cost of computing as deep networks tend to overfit on the dataset with limited samples per class. The model uses four parallel instances of convolution layer with varying kernel sizes. We experimented with different values for hyperparameters such as kernel number, kernel size, and finalized the following values based on the validation set performance:

- Kernel size:

$$f_1^h = 3, f_2^h = 6, f_3^h = 9, f_4^h = 12$$

- Number of kernels = 200, stride = 1.

3.2.1 Multi-Task Learning

In machine learning, multi-task learning is an old idea studied by Caruana (1997). A widely used technique to apply MTL is to train the main and auxiliary task jointly. In our work, the main task has text utterances which belong to either of the three classes, i.e. in favor, against and neutral whereas the proposed auxiliary task has two classes which comprise of neutral stance tweets and those which show a stance (in favor + against). The MTL framework allows the model parameter to be shared across tasks and enables the incorporation of a combined loss function with a shared underlying representation shown in Figure 1. Shared learning pushes the model to learn the feature representations that generalize well across tasks. The following loss function is comprised of loss of the main task and the auxiliary task. We use a lambda parameter to control the effect of loss of the auxiliary task on the total loss.

- Loss function:

$$L_{total} = L_{task1} + \lambda * L_{task2}$$

Here, λ is a tunable parameter which is optimized as part of the training process. We investigated the effectiveness of multi-task learning in an end-to-end neural network architecture for both the auxiliary task and the main task. We observed that the effect of task selection on model performance where it is validated that using auxiliary tasks improve the performance of the main task (Caruana, 1997).

Given suitable data, this approach is flexible enough to extend to other NLP tasks. It provides synergy between the two tasks, resulting in improved performance in comparison to individual tasks. The combined loss function pushes the model to learn general and complex features across multiple tasks rather than forcing the model to learn the features of a single task independently. This is a particularly interesting technique in NLP since data is scarce for many tasks and shared learning approach reduces the amount of training data needed.

4 Results

Model	Accuracy(%)
RBF Kernel SVM*	58.7
Random Forest*	54.7
Linear SVM*	56.6
CNN	**61.4**
CNN + MTL	**63.2**

Table 1: Detailed accuracies achieved on the benchmark dataset by different models. *RBF Kernel SVM, Random Forest, and Linear SVM accuracies are from (Swami et al., 2018)

The benchmark dataset that is published online by (Swami et al., 2018) is used for evaluating the effectiveness of machine translated input for our proposed architecture. It contains a total of 3545 annotated tweets where 1755 are labeled in favor, 647 as against and 1934 as neutral tweets. For the

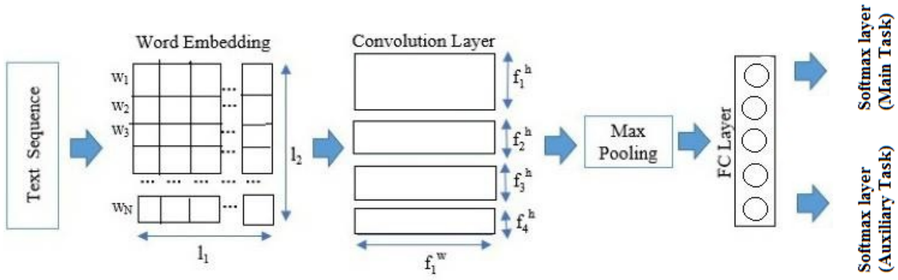

Figure 1: Proposed CNN - MTL architecture

Model	Accuracy(%)
CNN	66.7
CNN + MTL	71.3

Table 2: Comparison of accuracies for the auxiliary task

two tasks, we achieved an accuracy of 63.2%. We carried out 10-fold cross-validation for generating all our experimental results. Using all the features (Swami et al., 2018), the baseline systems: RBF kernel SVM, random forest, linear SVM presented an accuracy of 58.7%. Going forward, to the best of our knowledge, we are the first to experiment with deep learning architecture based on MTL for detecting stance in code-mixed data. The challenges in this task are the linguistic complexity and the lack of clean code-mixed data. And, preprocessing of code-mixed data will increase model performance.

In Table 1, we present the results of both the tasks with the proposed deep learning based architecture with translated data as input. We experimented with both continuous bag of words (CBOW) and skip-gram versions of word embeddings with CNN model and achieved similar results. The substantial accuracy obtained (63.2% for stance) shows more than 4.5% increment from values reported by (Swami et al., 2018). However, these values reflect that there is still a lot of room for improvement, justifying further efforts. We observed more than 4% overall accuracy improvement in the auxiliary task with the introduction of MTL as compared to the performance on the standalone CNN architecture. This indicates that training the main and the auxiliary task jointly can learn robust shared features which leads to improvement on both the main and auxiliary task.

5 Conclusion

We present MTL based deep learning approach for the problem of detecting user stance with respect to a particular topic: “Demonetisation”, on Twitter’s code-mixed Hindi-English data generated by bilingual users. The machine transliterated and translated corpus is given to the model. We empirically demonstrated the effectiveness of the proposed architecture. The proposed approach of jointly training the main and the auxiliary task proved to be the best-performing model so far for the code-mixed data, indicating that it is a promising new direction in the automated assessment of stance. An accuracy of 63.2% is achieved from our proposed deep learning model based on multi-task learning at detecting stance in code-mixed data which is an improvement of more than 4.5% overall accuracy when compared with current benchmark results.

6 Future Work

Our work provided insights regarding the benefits of training the main and the auxiliary task jointly for code-mixed data. There is a lot of room for improvement, and we hope to get a better understanding of how to improve the techniques for stance classification by primarily improving the corpus quality in our future work. Further, we will compare and contrast with different networks like LSTM, Attention-based architectures, etc. The results of our experiments are encouraging though since they show that it is possible to use classical methods for analyzing code-mixed texts. Furthermore, to address phrasal repetitions, short and simple constructions, non-grammatical words, more corpus without spelling errors need to be constructed as this can help other NLP tasks in multilingual societies.

References

Pranav Anand, Marilyn Walker, Rob Abbott, Jean E Fox Tree, Robeson Bowmani, and Michael Minor. 2011. Cats rule and dogs drool!: Classifying stance in online debate. In *Proceedings of the 2nd workshop on computational approaches to subjectivity and sentiment analysis*, pages 1–9. Association for Computational Linguistics.

Antonio Fernández Anta, Luis Núnez Chiroque, Philippe Morere, and Agustín Santos. 2013. Sentiment analysis and topic detection of spanish tweets: A comparative study of of nlp techniques. *Procesamiento del lenguaje natural*, 50:45–52.

Isabelle Augenstein and Anders Søgaard. 2017. Multi-task learning of keyphrase boundary classification. *arXiv preprint arXiv:1704.00514*.

Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.

Soravit Changpinyo, Hexiang Hu, and Fei Sha. 2018. Multi-task learning for sequence tagging: An empirical study. *arXiv preprint arXiv:1808.04151*.

Sarah Djemili, Julien Longhi, Claudia Marinica, Dimitris Kotzinos, and Georges-Elia Sarfati. 2014. What does twitter have to say about ideology? In *NLP 4 CMC: Natural Language Processing for Computer-Mediated Communication/Social Media-Pre-conference workshop at Konvens 2014*, volume 1, pages http–www. Universitätsverlag Hildesheim.

Adam Faulkner. 2014. Automated classification of stance in student essays: An approach using stance target information and the wikipedia link-based measure. In *FLAIRS Conference*.

Imene Guellil and Kamel Boukhalfa. 2015. Social big data mining: A survey focused on opinion mining and sentiments analysis. In *2015 12th International Symposium on Programming and Systems (ISPS)*, pages 1–10. IEEE.

Divam Gupta, Kushagra Singh, Soumen Chakrabarti, and Tanmoy Chakraborty. 2019. Multi-task learning for target-dependent sentiment classification. *arXiv preprint arXiv:1902.02930*.

Kazi Saidul Hasan and Vincent Ng. 2013. Stance classification of ideological debates: Data, models, features, and constraints. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1348–1356.

Peter Krejzl, Barbora Hourová, and Josef Steinberger. 2017. Stance detection in online discussions. *arXiv preprint arXiv:1701.00504*.

Man Lan, Jianxiang Wang, Yuanbin Wu, Zheng-Yu Niu, and Haifeng Wang. 2017. Multi-task attention-based neural networks for implicit discourse relationship representation and identification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1299–1308.

Navonil Majumder, Soujanya Poria, Haiyun Peng, Niyati Chhaya, Erik Cambria, and Alexander Gelbukh. 2019. Sentiment and sarcasm classification with multitask learning. *arXiv preprint arXiv:1901.08014*.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. Semeval-2016 task 6: Detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41.

Saif M Mohammad. 2016. Sentiment analysis: Detecting valence, emotions, and other affectual states from text. In *Emotion measurement*, pages 201–237. Elsevier.

Thai-Hoang Pham, Khai Mai, Nguyen Minh Trung, Nguyen Tuan Duc, Danushka Bolegala, Ryohei Sasano, and Satoshi Sekine. 2019. Multi-task learning with contextualized word representations for extented named entity recognition. *arXiv preprint arXiv:1902.10118*.

Ashwin Rajadesingan and Huan Liu. 2014. Identifying users with opposing opinions in twitter debates. In *International conference on social computing, behavioral-cultural modeling, and prediction*, pages 153–160. Springer.

Swapna Somasundaran and Janyce Wiebe. 2010. Recognizing stances in ideological on-line debates. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 116–124. Association for Computational Linguistics.

Sahil Swami, Ankush Khandelwal, Vinay Singh, Syed Sarfaraz Akhtar, and Manish Shrivastava. 2018. An english-hindi code-mixed corpus: Stance annotation and baseline system. *arXiv preprint arXiv:1805.11868*.

Marilyn A Walker, Pranav Anand, Robert Abbott, and Ricky Grant. 2012. Stance classification using dialogic properties of persuasion. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 592–596. Association for Computational Linguistics.