SSN_NLP@SardiStance : Stance Detection from Italian Tweets using RNN and Transformers

Kayalvizhi S
SSN College Of Engineering
kayalvizhis@ssn.edu.in

Thenmozhi D
SSN College Of Engineering
theni_d@ssn.edu.in

Aravindan Chandrabose
SSN College Of Engineering
aravindanc@ssn.edu.in

Abstract

Stance detection refers to the detection of one's opinion about the target from their statements. The aim of sardistance task is to classify the Italian tweets into classes of favor, against or no feeling towards the target. The task has two sub-tasks : in Task A, the classification has to be done by considering only the textual meaning whereas in Task B the tweets must be classified by considering the contextual information along with the textual meaning. We have presented our solution to detect the stance utilizing only the textual meaning (Task A) using encoder-decoder model and transformers. Among these two approaches, simple transformers have performed better than the encoder-decoder model with an average F1-score of 0.4707.

1 Introduction

Stance is the opinion of a person against or in favor of the target. In the sardistance task, the stance detection refers to the detection of stance from the Italian tweets collected from Sardines movement. The tweets imply the authors' standpoint towards the target. The aim of this task is to detect the stance of the author with the help of textual and contextual information about the tweets. The task has two sub-tasks in which the stance is detected using only textual information in one subtask while the other sub-task makes use of contextual meaning along with the textual meaning.

2 Related Work

Many approaches have been done to detect stance from the English text. Stance text are vectorized and then detected using Multi-layer Perceptron (MLP) (Riedel et al., 2017). Different methodologies like Support Vector Machine, Long Short Term Memory (LSTM) and Bi-directional LSTM (Augenstein et al., 2016) have also been used to detect stance. Recurrent Neural Network (RNN) (Yoon et al., 2019) and altering recurrent networks with different short connections pooling and attention layers have also been experimented in (Borges et al., 2019) to detect stance. Bi-directional Encoder Representation of Transformers (BERT) (Devlin et al., 2018) and Named Entity Recognition (NER) model (Küçük and Can, 2019) have also been used to detect stance. A large dataset has been collected from twitter and all the existing approaches have been discussed in (Conforti et al., 2020).

For other languages, a multilingual data set (Vamvas and Sennrich, 2020) have been taken, language is identified and then multi-lingual BERT model have been used to detect stance. Stance have been detected in Russian Language (Lozhnikov et al., 2018) by vectorizing using Tf-IDF and then classifying using different classifiers like Bagging, AdaBoost Boosting, Stochastic Gradient Descent classifier and Logistic Regression. Stance from different languages (Lai et al., 2020) like English, Italian, French, Spanish have been detected using different features extraction.

3 Task Description

The sardistance task (Cignarella et al., 2020) of Evalita (Basile et al., 2020) has two sub-tasks namely Task A - textual stance detection and Task B - contextual stance detection.
Both tasks are classification tasks that have three classes namely favor, against and none. In the first task, the system has to predict the class by using only the textual information from the tweets whereas in the second task it has to predict the label with the help of some additional information

like
Details of post : the number of re-tweets, replies, quotes
Details of user : the number of tweets, user bio's, user's number of friends and followers
Details of their social network : friends, replies, re-tweets, quotes' relation.
In both the tasks, there can be two submissions like constrained where we have to use only the dataset provided and unconstrained where we can use some additional data if required. Each team can submit two runs for both constrained and unconstrained runs.

3.1 Data set description

For Task A, the train.csv file was provided with three columns namely tweet_id,user_id and text label. For Task B, files namely tweet.csv, user.csv, friend.csv, quote.csv, reply.csv and re-tweet.csv are given to explain the contextual details about the post, user and social network. For both the tasks, the training set had about 2,132 instances and the test set had about 1,110 instances. In the training set, there are 1,028 instances in the against class, 587 favor instances and 515 neutral instances which is explained in Table 1. In the testing set, there are 742 against instances, 196 favor instances and 687 none instances.

4 Methodology

The stances were detected using an encoder-decoder model which is a recurrent neural network with different recurrent units and using transformers.

4.1 Data pre-processing

The data is pre-processed by removing the hash tags, '@' symbols, Unicode characters and punctuation.

4.2 Recurrent Neural Network

In this approach, the stance were detected using a encoder-decoder model (Luong et al., 2017) using Gated Recurrent unit(GRU) as its recurrent unit and Scaled Luong (Luong et al., 2015) as its attention mechanism. The model has two encoder-decoder layers along with the embedding layer that vectorizes the input and a loss layer that calculates the loss function. Recurrent Neural Network has been made use to detect the stance since it captures the contextual long-short term dependencies.

4.2.1 Encoder-Decoder Model

The encoder-decoder model is a Neural Machine Translation (NMT) model with sequential data model with Recurrent Neural Network (RNN). The Seq-to-Seq model differs in terms of type of recurrent unit, residual layers, depth, directionality and attention mechanism. The types of the recurrent unit are Long Short Term Memory(LSTM), Gated Recurrent Unit (GRU) and Google Neural Machine Translations. The depth is altered by changing the number of layers and the directionality is either uni-directionality or bi-directionality.The two types of attention mechanism are scaled luong (sl) and normed bahdanau (nb). The given training set is divided into development set and training set and the performance is measured using the development set which is shown in Table 2. The model was trained for about "10,000 steps", 6 epoch_step with "128 units", batch size of "128", dropout of "0.2" and learning rate of "0.1".

4.3 Transformers

In this approach, the stances were detected using simple transformers. Simple transformers are the wrapper of transformers. Transformers are mechanism that utilizes the attention mechanisms without using recurrent units. Bi-directional Encoder Representation of Transformers (BERT) is used to detect stance with the multilingual model and base model for the development set whose performance is given in Table 3. Multilingual Bert model (Devlin et al., 2018) of hugging face Pytorch transformers (Wolf et al., 2019) has been used to detect stance in our approach which was submitted as Run-1.

5 Results

Table 2 shows the different models evaluated based on the development set. From the table, the model with two layers of gated recurrent unit and scaled luong attention mechanism seems to perform better.

Table 4 shows the performance of various teams in this task of detecting stance. Twelve teams have participated in which one team have submitted both constrained and unconstrained runs which is denoted by the suffix "_u" in the table. Remaining all runs are constrained runs which are done only using the data set provided.

Data Distribution	against	favor	none	Total
Training set	1028	587	515	2132
Testing set	742	196	172	1110
Total instances	1770	783	687	3242

Table 1: Data distribution

Model name	Accuracy
2l_nb_gru	37.0
2l_sl_gru	**38.0**
3l_nb_gnmt	33.7
3l_sl_gnmt	33.7
4l_nb_gru	36.4
4l_sl_gru	35.7
3l_sl_gnmt_residual	37.5
3l_nb_gnmt_residual	37.5

Table 2: Performance of various models

Model	mcc	loss function
Bert- Multilingual	**0.167**	**1.098**
Bert - Base	0.141	1.150

Table 3: Performance of BERT models

The performance metrics used are class-wise prediction of precision, recall, F1-score and average F1-score. The ranking is done using an average F1-score which is shown in 4. The best performance in constrained run is 0.6801 whereas our approach of transformers (SSN_NLP run 1) has an average F1 score of 0.4707 and encoder-decoder model (SSN_NLP run 2) has an average score of 0.4473.

6 Conclusion

Italian tweets about the Sardines movement have been utilized to detect the opinion of the author towards the target. Different approaches have been made to detect the stance in the tweets by many other teams. We detected the stance using encoder-decoder model and simple transformers of multilingual Bert model in which transformers performed better than the encoder-decoder model with a F1-average score of 0.4707. The performance can further be improved by utilizing the additional dataset to train the model better to detect the stance in the tweets.

Acknowledgments

We would like to express our gratefulness towards DST-SERB funding agent and HPC laboratory of SSN College Of Engineering for providing space and resources required for this experiment.

References

Isabelle Augenstein, Tim Rocktäschel, Andreas Vlachos, and Kalina Bontcheva. 2016. Stance detection with bidirectional conditional encoding. *arXiv preprint arXiv:1606.05464*.

Valerio Basile, Danilo Croce, Maria Di Maro, and Lucia C. Passaro. 2020. Evalita 2020: Overview of the 7th evaluation campaign of natural language processing and speech tools for italian. In Valerio Basile, Danilo Croce, Maria Di Maro, and Lucia C. Passaro, editors, *Proceedings of Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2020)*. CEUR-WS.org.

Luís Borges, Bruno Martins, and Pável Calado. 2019. Combining similarity features and deep representation learning for stance detection in the context of checking fake news. *Journal of Data and Information Quality (JDIQ)*, 11(3):1–26.

Alessandra Teresa Cignarella, Mirko Lai, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020. SardiStance@EVALITA2020: Overview of the Task on Stance Detection in Italian Tweets. In Valerio Basile, Danilo Croce, Maria Di Maro, and Lucia C. Passaro, editors, *Proceedings of the 7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEUR-WS.org.

Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd, and Nigel Collier. 2020. Will-they-won't-they: A very large dataset for stance detection on twitter. *arXiv preprint arXiv:2005.00388*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding.

Dilek Küçük and Fazli Can. 2019. A tweet dataset annotated for named entity recognition and stance detection. *arXiv preprint arXiv:1901.04787*.

Team	F-average
SSN_NLP run 1 (transformers)	0.4707
SSN_NLP run 2 (encoder-decoder model)	0.4473
Team A - 1_u	0.6853
Team A - 1_c	0.6801
Team A - 2_c	0.6793
Team B - 1	0.6621
Team A - 2_u	0.6606
Team C - 1	0.6473
Team D - 1	0.6257
Team C - 2	0.6171
Team E	0.6067
Team B - 1	0.6004
Team D - 2	0.5886
Team F	0.5784
Team G - 1	0.5773
Team H	0.5749
Team I - 1	0.5595
Team I - 1	0.5329
Team J	0.4989
Team G - 2	0.4705
Team K	0.3637

Table 4: Results

Mirko Lai, Alessandra Teresa Cignarella, Delia Irazú Hernández Farías, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020. Multilingual stance detection in social media political debates. *Computer Speech & Language*, page 101075.

Nikita Lozhnikov, Leon Derczynski, and Manuel Mazzara. 2018. Stance prediction for russian: data and analysis. In *International Conference in Software Engineering for Defence Applications*, pages 176–186. Springer.

Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.

Minh-Thang Luong, Eugene Brevdo, and Rui Zhao. 2017. Neural machine translation (seq2seq) tutorial. *https://github.com/tensorflow/nmt*.

Benjamin Riedel, Isabelle Augenstein, Georgios P Spithourakis, and Sebastian Riedel. 2017. A simple but tough-to-beat baseline for the fake news challenge stance detection task. *arXiv preprint arXiv:1707.03264*.

Jannis Vamvas and Rico Sennrich. 2020. X-stance: A multilingual multi-target dataset for stance detection. *arXiv preprint arXiv:2003.08385*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rmi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

Seunghyun Yoon, Kunwoo Park, Joongbo Shin, Hongjun Lim, Seungpil Won, Meeyoung Cha, and Kyomin Jung. 2019. Detecting incongruity between news headline and body text via a deep hierarchical encoder. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 791–800.