

JU_NLP at SemEval-2016 Task 6: Detecting Stance in Tweets using Support Vector Machines

Braja Gopal Patra, Dipankar Das, and Sivaji Bandyopadhyay
Department of Computer Science and Engineering, Jadavpur University
Kolkata, India
{brajagopal.cse, dipankar.dipnil2005}@gmail.com
sivaji_cse_ju@yahoo.com

Abstract

We describe the system submitted to the SemEval-2016 for detecting stance in tweets (Task 6, Subtask A). One of the main goals of stance detection is to automatically determine the stance of a tweet towards a specific target as ‘FAVOR’, ‘AGAINST’, or ‘NONE’. We developed a supervised system using Support Vector Machines to identify the stance by analyzing various lexical and semantic features. The average F1 score achieved by our system is 60.60.

1 Introduction

Recent research in the areas of opinion mining and/or sentiment analysis on natural language texts is gaining importance due to various academic and commercial perspectives. One of the main reasons is that a vital amount of information can be obtained from text data that are available in the internet in forms of news, reviews, blogs, chats, and tweets.

Several experiments have been attempted in the field of sentiment analysis or opinion mining on user generated content or social media data till date (Patra et al., 2015). One of the main goals in such tasks is to assign polarities (positive or negative) to a piece of text. Identification of the writer’s attitude towards a specific target, known as stance detection, is a relevant and challenging topic which has unduly been averted by most researchers (Mohammad et al., 2016). Lately it has become conventional that people express stance explicitly or implicitly in various microblogging sites.

Stance detection is the task of automatically determining from text whether the author is in favor of

the given target, against the given target, or whether neither inference is likely.¹ The stance detection can often bring complementary information to sentiment analysis, because we often care about the author’s evaluative outlook towards specific targets and propositions rather than simply considering whether the speaker is angry or happy (Mohammad et al., 2016). Stance detection even becomes more difficult when it is performed on the short texts like tweets. The latter being an important sub-field of sentiment analysis/opinion mining. Automatic stance detection can be used in several applications such as information retrieval, text summarization, and textual entailment.

More recent approaches to stance detection have been performed using linguistic rules on *online debate dataset* (Somasundaran and Wiebe, 2009; Hasan and Ng, 2013; Sridhar et al., 2014) and *NPOV Corpus* from Wikipedia (Recasens et al., 2013). To the best of our knowledge, not much computational attempts have been performed on tweets for stance detection. It is also observed that both the supervised and unsupervised machine learning algorithms have been implemented for identifying the stance and several features like n-gram, frame semantic features, and dependency were used in stance detection tasks (Somasundaran and Wiebe, 2009; Anand et al., 2011; Sridhar et al., 2014).

We participated in the SemEval 2016-Task 6: Detecting Stance in Tweets.¹ The main goal of this task is to identify stance present in a tweet towards a specific target. For example, **Target** = ‘Hillary Clinton’, **Tweet** = ‘A Black President, Healthcare 4

¹<http://alt.qcri.org/semeval2016/task6/>

all, Marriage Equality...what's next?...a Woman President?!Damn RIGHT! #SCOTUS #MarriageEquality, **Stance** = 'FAVOR'. In the above example, the tweet is in favor of the target "Hillary Clinton". In the test dataset, the targets of the corresponding tweets are given and we have to find out whether the stance towards the target is in 'FAVOR' or 'AGAINST' or 'NONE'.

There are two subtasks, in subtask A, we have to identify the stance towards a specific target whereas in subtask B, we have to identify the stance towards only one target "Donald Trump" using a large set of unlabeled tweets associated with the target. We have only participated in subtask A. Several lexical and semantic features are used to identify the stance towards a target. Support Vector Machines (SVM) is used for the classification purpose.

2 Dataset and Evaluation

The organizers provided the trial, training, and test dataset of 100, 2814, and 1249 tweets, respectively for detecting stance towards five targets namely "Atheism", "Climate Change is a Real Concern", "Feminist Movement", "Hillary Clinton", and "Legalization of Abortion". All tweets in the dataset are annotated with stance ('FAVOR' or 'AGAINST' or 'NONE') towards the above mentioned targets. The stance 'NONE' means the tweet is not in 'FAVOR' or 'AGAINST' of the target.

Systems were evaluated using the official metric calculated macro-average of two labels only, i.e. **F1 score** = $(F\text{-SCORE}_{FAVOR} + F\text{-SCORE}_{AGAINST})/2$.

3 System Framework

3.1 Features

In order to achieve good accuracy using machine learning algorithms, we need to have a good set of features. Based on the genre of tweet, we identified the following features as listed below.

3.1.1 Target Specific Words

Initially, an n-gram model was created to identify the target's presence in the tweet. But, we observed that the dimension as well as the sparsity of feature vectors were increased enormously. Thus, we created individual topic bags related to

	Topic Bags	Word Bags	
		Favor	Against
A	311	69	210
C	129	63	42
F	95	72	49
H	73	40	31
L	101	44	59

Table 1: Statistics of the topic, favor, and against bags for five targets (A: Atheism, C: Climate Change is a Real Concern, F: Feminist Movement, H: Hillary Clinton, L: Legalization of Abortion)

each of the targets. Seed words related to the targets were populated using RitaWordNet.² Some handcrafted rules were implemented to find names in words and hashtags. For example, "Hillary" is present in "#HillaryClinton", then we considered "#HillaryClinton" in the topic bag "Hillary". Further, we checked the topic bags manually and removed some unrelated words collected using the WordNet. The detailed word level statistics of the topic bags are given in Table 1.

At first, we checked that whether a tweet contains any word from the corresponding target related topic bag or not. If no word is present, then we tagged that tweet as 'NONE' and no further processing is performed on that tweet. Further, this topic bags are used along with the dependency information.

3.1.2 Sentiment Words

We have used three lexicons namely SentiWordNet (Baccianella et al., 2010), NRC Emotion Lexicon (Mohammad and Turney, 2013), and NRC Hashtag Emotion Lexicon (Mohammad, 2012) for our experiments along with manually created lexicons for each of the targets. Moreover, two bags were created for each of the targets namely 'favor' and 'against' bags. Hashtags are also included in the above lexicons. The word level statistics of the 'favor' and 'against' bags are given in Table 1. We used the frequency of the sentiment words matched with the above sentiment or emotion lexicons as features. Again, we used these lexicons in the later experiments also.

²<https://rednoise.org/rita/reference/RiWordNet.html>

Parse

```
(ROOT
 (S
  (S
   (NP (NN @HillaryClinton) (NNS i))
   (VP (VBP am)
        (ADJP (JJ portuguese))))
  (CC and)
  (S
   (NP (FW i))
   (VP (VBP support)
        (NP (PRP$ your) (NN campaign))))
  (. !)))

(ROOT
 (INTJ (JJ Best) (NN luck) (. !)))
```

Universal dependencies

```
compound(i-2, @HillaryClinton-1)
nsubj(portuguese-4, i-2)
cop(portuguese-4, am-3)
root(ROOT-0, portuguese-4)
cc(portuguese-4, and-5)
nsubj(support-7, i-6)
conj(portuguese-4, support-7)
nmod:poss(campaign-9, your-8)
dobj(support-7, campaign-9)

amod(luck-2, Best-1)
root(ROOT-0, luck-2)
```

Figure 1: Dependency relations from Stanford Parser

3.1.3 Dependency Information

It is found in the literature that the dependency relations act as useful feature in sentiment analysis (Patra et al., 2014a; Patra et al., 2014b). Thus, we used the Stanford Parser³ to get the dependency relations. We searched for the word pairs in dependency relations that consist of two component words, one of which is present in either ‘favor’ or ‘against’ bag, whereas the other one is present in SentiWordNet.

In the above Figure 1, we found a relation “**dobj(support-7, campaign-9)**”. The word ‘campaign’ is present in the ‘favor’ bag for the target “*Hillary Clinton*” and the word ‘support’ is present in SentiWordNet as positive. This relation is considered as favor_positive type. Similarly, three other types namely favor_negative, against_positive, and against_negative were considered. The counts of each type are used as features.

Again, we identified the simple sentences based on the symbol “(S” or “(ROOT” in Figure 1. For

³<http://stanfordnlp.github.io/CoreNLP/>

each of the simple sentences, we removed the stop-words (except the negation words) and identified the sentiment words using SentiWordNet. We applied some handcrafted rules given below. 1. Not + NEG = POS, 2. Not + POS = NEG, 3. POS + POS = POS, 4. NEG + NEG = NEG. If there are multiple positive words present in a simple sentence, we tagged it as positive. If there are only two sentiment words, one with positive and other as negative, we tagged the sentence as negative. It is observed that the frequency of the negative instances is much higher than the positive instances in the training data. Finally, we counted the number of positive and negative instances in a tweet and used as features.

3.2 Classification Framework

We used LibSVM, an variant of SVMs implemented in the Weka.⁴ We performed the 10-fold cross validation on the training and trial dataset for all the targets using the above features. But, the system performance was observed as poor and the system achieves the average F1 score of 0.43. Again, we performed the 10-fold cross validation on the training and trial dataset for all the targets, separately. We calculated the average F1 score for each of the five targets and these are 0.46, 0.40, 0.48, 0.52, and 0.49 for “Atheism”, “Climate Change is a Real Concern”, “Feminist Movement”, “Hillary Clinton”, and “Legalization of Abortion”, respectively. This yields an average F1 score of 0.47 for all the targets, which is higher than the F1 score achieved by the previous system. This motivates us to develop our final systems separately for each of the targets.

3.2.1 Results

The LibSVM based system achieved the maximum F-Score of 0.4668 and 0.7452 for the ‘FAVOR’ and ‘AGAINST’ classes, respectively as shown in Table 2. The system achieved the maximum average F1 score of 0.6060. Whereas, the team MITRE has achieved the first position with the maximum average F1 score of 0.6782.

The confusion matrix for our system is shown in Table 3. From the matrix, we observed that the system is biased towards the ‘AGAINST’ stance (except “Climate Change is a Real Concern” target). The main reason may be the biasness in the

⁴<http://www.cs.waikato.ac.nz/ml/weka/>

	Precision	Recall	F-Score
FAVOR	0.6687	0.3586	0.4668
AGAINST	0.7227	0.7692	0.7452

Table 2: System performance

		Predicted			
		A	F	N	
Actual	Atheism	A	124	2	34
		F	25	0	7
		N	9	0	19
	Climate Change is a Real Concern	A	1	7	3
		F	21	92	10
		N	1	10	24
	Feminist Movement	A	145	12	26
		F	45	6	7
		N	9	2	33
	Hillary Clinton	A	135	15	22
		F	33	9	3
		N	21	1	56
	Legalization of Abortion	A	145	4	40
		F	35	2	9
		N	12	1	32

Table 3: Confusion Matrix of our System (A: Against, F: Favor, N: None)

training dataset as the number of instances for the ‘AGAINST’ stance is maximum for all the targets (except “Climate Change is a Real Concern” target). Another reason may be the less number of training instances as we developed our system for each of the targets, separately.

4 Conclusion and Future Work

We presented a system for identifying the stance in tweets using dependency and semantic features. The maximum average F1 score of 0.6060 is achieved by the system using SVM classifier. The task stance detection on social media data is helpful for real life applications like political campaign and opinion poll.

In near future, we plan to use the TweeboParser,⁵ as it works well for tweets as compared to the Stanford Parser. Another immediate goal is to increase the size of topic bags and sentiment lexicons related to each of the targets. We are also planning to use unsupervised approach and several machine learning

⁵<http://www.cs.cmu.edu/ark/TweetNLP/>

algorithms for classifying the stance in near future.

Acknowledgments

The present work is supported by a grant from the project “CLIA System Phase II” funded by Department of Electronics and Information Technology (DeitY), Ministry of Communications and Information Technology (MCIT), Government of India. The first author is supported by “Visvesvaraya Ph.D. Fellowship” funded by DeitY, Government of India.

References

- Pranav Anand, Marilyn Walker, Rob Abbott, Jean E. F. Tree, Robeson Bowmani, and Michael Minor. 2011. Cats rule and dogs drool!: Classifying stance in on-line debate. In *Proceedings of the 2nd workshop on computational approaches to subjectivity and sentiment analysis*, pages 1–9.
- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of 7th Language Resources and Evaluation Conference*, volume 10, pages 2200–2204.
- Kazi S. Hasan and Vincent Ng. 2013. Stance classification of ideological debates: Data, models, features, and constraints. In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, pages 1348–1356.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 Task 6: Detecting stance in tweets. In *Proceedings of the International Workshop on Semantic Evaluation*, SemEval’16, San Diego, California.
- Saif M. Mohammad. 2012. #emotional tweets. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics*, pages 246–255, Montréal, Canada, 7-8 June. Association for Computational Linguistics.
- Braja G. Patra, Soumik Mandal, Dipankar Das, and Sivaji Bandyopadhyay. 2014a. Ju_cse: A conditional random field (crf) based approach to aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 370–374.
- Braja G. Patra, Niloy Mukherjee, Arijit Das, Soumik Mandal, Dipankar Das, and Sivaji Bandyopadhyay.

- 2014b. Identifying aspects and analyzing their sentiments from reviews. In *Proceedings of 13th Mexican International Conference on Artificial Intelligence (MICA)*, pages 9–15. IEEE.
- Braja G. Patra, Dipankar Das, Amitava Das, and Rajendra Prasath. 2015. Shared task on sentiment analysis in indian languages (sail) tweets-an overview. In *Proceedings of Mining Intelligence and Knowledge Exploration*, pages 650–655. Springer.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic models for analyzing and detecting biased language. In *Proceedings of the 51th Annual Meeting of the ACL*, pages 1650–1659.
- Swapna Somasundaran and Janyce Wiebe. 2009. Recognizing stances in online debates. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing*, pages 226–234. Association for Computational Linguistics.
- Dhanya Sridhar, Lise Getoor, and Marilyn Walker. 2014. Collective stance classification of posts in online debate forums. pages 109–117.