

Binarizer at SemEval-2018 Task 3: Parsing dependency and deep learning for irony detection

Nishant Nikhil

IIT Kharagpur

Kharagpur, India

nishantnikhil@iitkgp.ac.in

Muktabh Mayank Srivastava

ParallelDots, Inc.

muktabh@paralleldots.com

Abstract

In this paper, we describe the system submitted for the SemEval 2018 Task 3 (Irony detection in English tweets) Subtask A by the team Binarizer. Irony detection is a key task for many natural language processing works. Our method treats ironical tweets to consist of smaller parts containing different emotions. We break down tweets into separate phrases using a dependency parser. We then embed those phrases using an LSTM-based neural network model which is pre-trained to predict emoticons for tweets. Finally, we train a fully-connected network to achieve classification.

1 Introduction

The micro-blogging site Twitter has created an abundance of data about opinions and sentiments regarding almost every aspect of daily life. A deeper study of the public opinion can be obtained by applying natural language processing techniques on this data. However, the performance of these NLP models is detrimentally affected by irony (Pozzi et al., 2016). As per the Oxford English Dictionary, irony is the expression of one’s meaning by using language that normally signifies the opposite, typically for humorous or emphatic effect. This deviation between what is said and what is intended makes irony hard to detect. Being a platform where users are free to communicate and express themselves colloquially, Twitter generates considerable data injected with irony. Studying this would provide us with a better sentiment analysis of these tweets.

Prior work on irony detection includes the use of unigrams and emoticons (González-Ibáñez et al., 2011; Carvalho et al., 2009; Barbieri et al., 2014). Maynard and Greenwood (2014) describe an unsupervised pattern mining approach where the sentiment of the hashtag in the tweet is proposed to be a key indicator of sarcasm. If the sentiment of the tweet does not match the sentiment

of the hashtag, it is predicted to be sarcastic. Riloff et al. (2013) illustrates a semi-supervised approach where rule-based classifiers are used to look for negative situation phrases and positive verbs in a sentence. Tsur et al. (2010) build pattern-based features that detect the presence of discriminative patterns as extracted from a large sarcasm-labelled corpus. N-gram-based approaches have also been used (Davidov et al., 2010; Ptáček et al., 2014; Joshi et al., 2015) with sentiment features. Joshi et al. (2017) use similarity between word embeddings as feature and Poria et al. (2016) use convolutional neural networks to extract sentiment, emotion and personality features for sarcasm detection.

SemEval-2018 Task 3 (the 12th workshop on semantic evaluation) specifies two subtasks in relation to irony detection in English tweets (Van Hee et al., 2018). In subtask A the goal was to train a binary classifier that detects whether a given tweet is ironic or not. Subtask B was a multi-class classification problem where four labels were specified to describe the nature of irony (verbal irony by means of a polarity contrast, situational irony, other verbal irony, and non-ironic). The goal was to assign one of the four labels to each tweet.

We propose a new method which considers ironical tweets to be collections of smaller parts containing different emotions. We break down tweets into these collections using a dependency parser and embed them using DeepMojji (Felbo et al., 2017) which is pre-trained to predict emoticons for tweets. Finally we train a classifier to detect irony. The paper is organized as follows: We discuss our methods in section 2. Section 3 contains the details about the experiments and the training data. In Section 4 we discuss the results and Section 5 concludes the paper with closing remarks.

2 Method

In order to identify the chunks of various emotions in an ironic tweet, we split the tweets into phrases using a dependency parser. We use Tweepoparser (Kong et al., 2014), which is a dependency parser for English tweets. The parser is trained on a subset of a labelled corpus for 929 tweets (12,318 tokens) drawn from the POS-tagged tweet corpus of Owoputi et al. (2013), Tweepbank. TweepoParser predicts the syntactic structure of the tweet represented by unlabelled dependencies. Tweets contain multiple sentences or fragments called “utterances” each with their own syntactic root disconnected from the others. Since a tweet often contains more than one utterance, the output of TweepoParser will often be a multi-rooted graph over the tweet. Also, many elements in tweets have no syntactic function. These include, in many cases, hashtags, URLs, and emoticons. TweepoParser attempts to exclude these tokens from the parse tree. For our purpose, we club the words arising from the same root to create a phrase. Multiple roots would create multiple phrases. As we can see from Figure 1, these phrases can convey the different sentiments attached to the different subjects of the tweet.

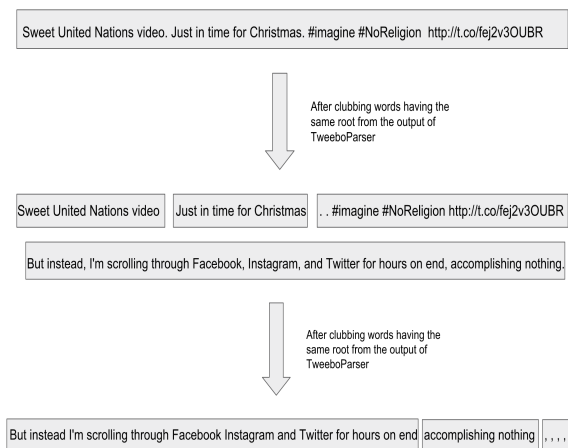


Figure 1: Parser results

After extracting a set of phrases for the sentence, we embed the phrases into vectors. We used the DeepMoji (Felbo et al., 2017) model, which is trained on 1.2 billion tweets with emojis to understand how language is used to express emotions. It encodes the provided phrase into a 2,304-dimensional feature vector. Under the hood, DeepMoji model projects each word into a 256-dimensional vector space followed by a hyperbolic

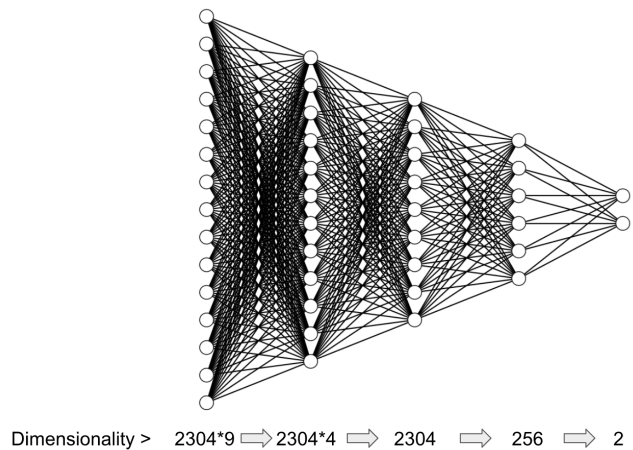


Figure 2: Neural network architecture

tangent activation function. After this, two bidirectional LSTMs with 1,024 hidden units each are used to capture the context of each word. Finally, the model uses skip connections from each layer to an attention layer and hence the attention layer outputs a 2,304 (256+1,024+1,024) dimensional vector. Now this 2,304-dimensional output is connected to a softmax layer for classification. We did not use the final softmax layer but took the 2,304-dimensional vector for each phrase. As the model was trained for prediction of emoticons, this feature vector contains information about the semantic and sentimental content of the phrases. To make the predictions we need to account for the sentiment behind every utterance. To this end, we concatenate these vectors and pass the resulting concatenated vector through a fully-connected network as described in Figure 2.

Tweets can have a varying number of roots, which implies that they split into a varying number of phrases. Our model considers a maximum of nine roots. A tweet with an excess of nine roots is truncated suitably. On the other hand, a tweet with less than nine roots is zero-padded. We have described the complete process flowchart in Figure 3.

3 Experiments

For subtask A, we were provided with a dataset consisting of tweets along with a binary class (0 or 1) which indicates whether this tweet is ironic or not (0 for non-ironic tweets and 1 for ironic tweets). The data was collected from Twitter API by querying tweets using the hashtags #irony, #sarcasm and #not, with subsequent manual an-

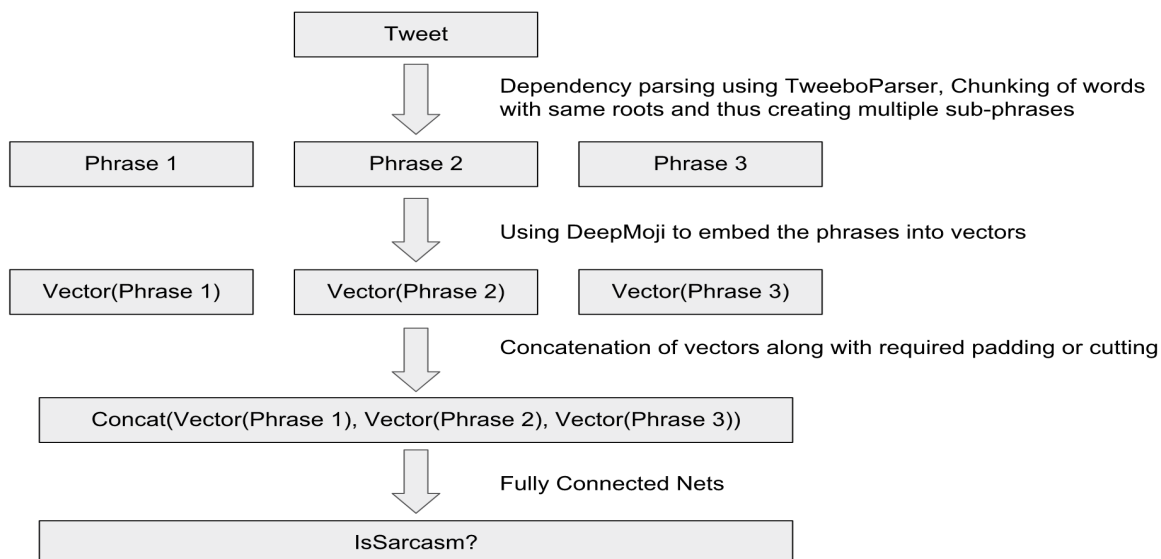


Figure 3: Process Flowchart

notation to remove noise. 3,833 tweets for training and 784 tweets for testing were provided. The evaluation was done by using accuracy, precision, recall and F1 score.

$$Accuracy : \frac{\text{number of correctly predicted instances}}{\text{total number of instances}}$$

$$Precision : \frac{\text{number of correctly predicted instances}}{\text{number of predicted labels}}$$

$$Recall : \frac{\text{number of correctly predicted instances}}{\text{number of labels in the gold standard}}$$

$$F1 \text{ score} : \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$$

We used the pipeline described in Figure 3. The final step of the process used a fully connected neural network with four layers. The input layer of the FC network has a dimension of 20,736 (2,304*9), the second layer has a dimension of 9,216 (2,304*4), the third layer has a dimension of 2,304 and the fourth layer has 256 dimensions. The final layer has 2 dimensions, with one for each class. This is depicted in Figure 2. We used the hyperbolic tangent activation function in all of the layers, and stochastic gradient descent with a learning rate of 0.01 and a momentum of 0.5.

Two models were then devised. The difference in these models lies in the input supplied to the FC network. In the first model, this input is the concatenation of the vectors obtained by embedding phrases. In the second model, the input is the

Method	Accuracy	Precision	Recall	F1 score
Winning team	0.7347	0.6304	0.8006	0.7054
Our System β_1	0.6659	0.5527	0.6471	0.5962
Our System β_2	0.6390	0.5198	0.6941	0.5944
Our System α	0.6951	0.6197	0.5176	0.5641

Table 1: Results SemEval Task 3A

concatenation of the input in the first model along with a 2,304-dimensional vector representing the embedding of the tweet as a whole. The results we get from various experiments on these models are shown in Table 1.

System α is the first model. The best F1 score for this model was achieved after four epochs, as shown in Table 1. System β_1 and β_2 are the second model running for five and four epochs respectively.

4 Results and Discussions

We participated only in the shared task 3A as the team Binarizer. We came ninth as per accuracy and seventeenth as per F1 score among the forty-three participating systems. Due to a glitch on

our side during submission the results are based on 446 out of 784 instances in the test data.

The models perform better than the baseline system as per the competition leaderboard. This reinforces the notion that separate phrases in a tweet carry information required for irony detection. System α has greater precision whereas System β has higher recall. So an application which demands urgent detection of ironic tweets would profit more from System β . This demonstrates that the sentiment information of the context provided from the whole tweet is also important.

5 Conclusion and Future works

We have shown how using the sentiments of different segments of tweets can enable irony detection. From the results of our experiments, we conclude that the segments have sufficient sentiment information in them for the identification of irony. In future research, we aim to improve the algorithm for parsing these chunks by replacing the dependency parser. Also, more experimentation can be performed for the last part of the pipeline. As the phrases from the tweets are sequences themselves, we can apply sequence modelling with LSTMs or CNNs.

References

- Francesco Barbieri, Horacio Saggion, and Francesco Ronzano. 2014. Modelling sarcasm in twitter, a novel approach. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 50–58.
- Paula Carvalho, Luís Sarmiento, Mário J Silva, and Eugénio De Oliveira. 2009. Clues for detecting irony in user-generated contents: oh...!! it's so easy;- . In *Proceedings of the 1st international CIKM workshop on Topic-sentiment analysis for mass opinion*, pages 53–56. ACM.
- Dmitry Davidov, Oren Tsur, and Ari Rappoport. 2010. Semi-supervised recognition of sarcastic sentences in twitter and amazon. In *Proceedings of the fourteenth conference on computational natural language learning*, pages 107–116. Association for Computational Linguistics.
- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1615–1625.
- Roberto González-Ibáñez, Smaranda Muresan, and Nina Wacholder. 2011. Identifying sarcasm in twitter: a closer look. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers-Volume 2*, pages 581–586. Association for Computational Linguistics.
- Aditya Joshi, Pushpak Bhattacharyya, and Mark J. Carman. 2017. Automatic sarcasm detection: A survey. *ACM Computing Surveys*, 50(5):73:1–73:22.
- Aditya Joshi, Vinita Sharma, and Pushpak Bhattacharyya. 2015. Harnessing context incongruity for sarcasm detection. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, volume 2, pages 757–762.
- Lingpeng Kong, Nathan Schneider, Swabha Swayamdipta, Archana Bhatia, Chris Dyer, and Noah A Smith. 2014. A dependency parser for tweets. In *Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*, pages 1001–1012, Doha, Qatar.
- Diana Maynard and Mark Greenwood. 2014. Who cares about Sarcastic Tweets? Investigating the Impact of Sarcasm on Sentiment Analysis. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4238–4243, Reykjavik, Iceland. European Language Resources Association.
- Olutobi Owoputi, Brendan O'Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A Smith. 2013. Improved part-of-speech tagging for online conversational text with word clusters. In *The 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2013)*, pages 380–390.
- Soujanya Poria, Erik Cambria, and Alexander Gelbukh. 2016. Aspect extraction for opinion mining with a deep convolutional neural network. *Knowledge-Based Systems*, 108:42–49.
- Federico Alberto Pozzi, Elisabetta Fersini, Enza Messina, and Bing Liu. 2016. *Sentiment analysis in social networks*. Morgan Kaufmann Publishers Inc., San Francisco, CA.
- Tomáš Ptáček, Ivan Habernal, and Jun Hong. 2014. Sarcasm detection on czech and english twitter. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 213–223.
- Ellen Riloff, Ashequl Qadir, Prafulla Surve, Lalin-dra De Silva, Nathan Gilbert, and Ruihong Huang. 2013. Sarcasm as contrast between a positive sentiment and negative situation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 704–714.

Oren Tsur, Dmitry Davidov, and Ari Rappoport. 2010. Icwsm-a great catchy name: Semi-supervised recognition of sarcastic sentences in online product reviews. In *Proceedings of International AAAI Conference on Web and Social Media*, pages 162–169.

Cynthia Van Hee, Els Lefever, and Vronique Hoste. 2018. Semeval-2018 task 3: Irony detection in english tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation (SemEval-2018)*, New Orleans, LA, USA.