ntuer at SemEval-2019 Task 3 : Emotion classification with word and sentence representations in RCNN

Zhong, Peixiang; Miao, Chunyan

2019

Zhong, P., & Miao, C. (2019). ntuer at SemEval-2019 Task 3 : Emotion classification with word and sentence representations in RCNN. Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval-2016).

<https://hdl.handle.net/10356/106025>

© 2019 Association for Computational Linguistics (ACL). All rights reserved. This paper was published in Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval-2016) and is made available with permission of Association for Computational Linguistics (ACL).

Downloaded on 26 Apr 2023 20:57:39 SGT

ntuer at SemEval-2019 Task 3: Emotion Classification with Word and Sentence Representations in RCNN

Peixiang Zhong, Chunyan Miao

School of Computer Science and Engineering

Nanyang Technological University Singapore

peixiang001@e.ntu.edu.sg, ascymiao@ntu.edu.sg

Abstract

In this paper we present our model on the task of emotion detection in textual conversations in SemEval-2019. Our model extends the Recurrent Convolutional Neural Network (RCNN) by using external fine-tuned word representations and DeepMoji sentence representations. We also explored several other competitive pre-trained word and sentence representations including ELMo, BERT and InferSent but found inferior performance. In addition, we conducted extensive sensitivity analysis, which empirically shows that our model is relatively robust to hyper-parameters. Our model requires no handcrafted features or emotion lexicons but achieved good performance with a micro-F1 score of 0.7463.

1 Introduction

Emotions are psychological and physiological states generated in humans in reaction to internal or external events. Messages in human conversations inherently convey emotions. With the rise of social media platforms such as Twitter, as well as chatbots such as Amazon Alexa, there is an emerging need for machines to understand human emotions in conversations, which has a wide range of applications such as opinion analysis in customer support (Devillers et al., 2002) and providing emotion-aware responses (Zhong et al., 2019). SemEval-2019 Task 3: EmoContext (Chatterjee et al., 2019b) is designed to promote research in this task.

This task is to detect emotions in textual conversations. Each conversation is composed of three turns of utterances and the objective is to detect the emotion of the last utterance given the first two utterances as the context. The emotions in this classification task include happy, sad, angry and others, adapted from the well-known Ekman’s six basic emotions: anger, disgust, fear, happiness,

sadness, and surprise (Ekman, 1992). The evaluation criteria is micro-averaged F1 score since the data is extremely unbalanced, as shown in Table 1.

In recent years, pre-trained word and sentence representations achieved very competitive performance in many NLP tasks, e.g., fine-tuned word embeddings using distant training (Cliche, 2017) and tweet sentence representations DeepMoji (Felbo et al., 2017) on sentiment analysis, and contextualized word representations BERT (Devlin et al., 2018) on 11 NLP tasks. Motivated by these successes, in this task we explored different word and sentence representations. We then fed these representations into a Recurrent Convolutional Neural Network (RCNN) (Lai et al., 2015) for classification. RCNN includes a Long short-term memory (LSTM) network (Hochreiter and Schmidhuber, 1997) to capture word ordering information and a max-pooling layer (Scherer et al., 2010) to learn discriminative features. We also experimented LSTM and CNN in our preliminary analysis but achieved worse performance as compared to RCNN. Our final system adopted fine-tuned word embeddings and DeepMoji as our choices of word and sentence representations, respectively, due to their superior performance on the validation dataset. The code is publicly available at Github¹.

2 Related Work

Emotion detection in textual conversations is an under-explored research task. The majority of existing works focused on the multi-modality settings (Devillers et al., 2002; Hazarika et al., 2018; Majumder et al., 2019). Chatterjee et al. (2019a) is one of the early works on the textual modality that first collected the dataset used in this task and then

¹<https://github.com/zhongpeixiang/SemEval2019-Task3-EmotionDetection>

Dataset Split	Size	#Happy	#Sad	#Angry	#Others	Average Utterance Length
Train	30160	4243	5463	5506	14948	5.22
Val	2755	142	125	150	2338	5.05
Test	5509	284	250	298	4677	5.05

Table 1: Total number of conversations and their distributions over each emotion class for each dataset split. Average number of tokens per utterance for each dataset split are also reported.

proposed an LSTM model with both semantic and sentiment embeddings to classify emotions. This task is also closely related to sentiment analysis (Pang et al., 2008) where the opinions of a piece of text is to be identified. One major difference between them is that this task detects emotions only on the last portion of a piece of text and the rest is treated as context.

Our model leverages pre-trained word and sentence representations. There is a research trend on word and sentence embeddings after the invention of Word2Vec (Mikolov et al., 2013). Cliche (2017) fine-tuned word embeddings using CNN-based sentiment classification model and distant training (Go et al., 2009). Peters et al. (2018) proposed a contextualized word embedding named ELMo to incorporate context information and solve the polysemy issues in conventional word embeddings. Devlin et al. (2018) proposed another contextualized word embedding named BERT by extending the context to both directions and training on the masked language modelling task. Kiros et al. (2015) proposed a sentence-level representation named SkipThought, which shares similar ideas to Word2Vec but operates on sentence level. Conneau et al. (2017) proposed InferSent by learning sentence representations on natural language inference tasks. Felbo et al. (2017) proposed DeepMoji by learning tweet sentence representations in the emoji classification task using 1246 million tweets and distant training.

Our RCNN model is closely related to neural network based sentiment analysis models. Two of the most popular models are LSTMs and CNNs. LSTM-based models can capture the word ordering information and have achieved the state-of-the-art performance on many sentiment analysis datasets (Gray et al., 2017; Liu et al., 2018; Howard and Ruder, 2018). CNN-based models can capture local dependencies, discriminative features, and are parallelizable for efficient computation (Kim, 2014; Johnson and Zhang, 2017).

3 System Description

In this section we describe our data preprocessing procedures and illustrate how we leverage pre-trained word and sentence representations in our RCNN model. The overall architecture is depicted in Figure 1.

3.1 Data Preprocessing

We concatenated three utterances as one sentence, separated by EOS tokens. We used the tokenizer from Spacy² for tokenization. We removed training sentences that have more than 75 tokens. We removed duplicate punctuations and spaces. We kept all remaining tokens in the training dataset as the vocabulary.

3.2 Pre-trained Word Representation

We fine-tuned the word embeddings obtained from (Baziotis et al., 2017), which has an embedding size of 100 and is pre-trained on 330M English Twitter messages using Glove (Pennington et al., 2014). The fine-tuning is conducted on the binary sentiment classification task using the basic CNN model (Kim, 2014) on 1.6 million tweets (Go et al., 2009). These tweets are labelled with positive and negative sentiments. Fine-tuning on these tweets introduces sentiment-discriminative features to word embeddings (Cliche, 2017). The CNN model has kernel sizes of 1, 2, and 3. Each kernel size has 300 filters. During fine-tuning, the embedding layer is first frozen for one epoch and then unfrozen for another three epochs.

3.3 Pre-trained Sentence Representation

We adopted DeepMoji (Felbo et al., 2017) as the sentence representations in our model. Each sentence will be encoded into a vector of size 2304. DeepMoji is trained on the 64-class emoji classification task using 1246 million tweets. Since emoji reflects emotions and sentiments, DeepMoji is an ideal model to provide emotion-discriminative sentence representations. We also

²<https://spacy.io/>

Figure 1: Overall architecture of our proposed model.

explored InferSent (Conneau et al., 2017), another sentence representation model with competitive performance on sentence classification and information retrieval tasks (Perone et al., 2018).

3.4 RCNN

As shown in Figure 1, we fed word and sentence representations into a RCNN model. The RCNN model mainly comprises of a two-layer Bi-directional LSTM (BiLSTM), a linear transformation layer and a max-pooling layer. At the embedding layer, each sentence is transformed to a sequence of word embeddings W_i of size 100 using our pre-trained word representations, where $i = 1, 2, \dots, n$, and n is the number of tokens in the concatenated utterance. The BiLSTM encodes these word embeddings into hidden states h_i^f, h_i^b in both forward and backward directions, respectively, where each direction has a hidden size of 200. The hidden states in both directions are concatenated together, along with the word representations W_i to form a vector of size 500. A linear transformation is then applied to project the resulted vector into a vector of size 200, followed by a max-pooling layer to extract discriminative sentence features. Finally, the DeepMojji sentence representation is concatenated with the pooled vector to form a final sentence representation of size 2504, followed by a softmax layer for classification.

3.5 Training

We train our model on the training dataset and fine-tune on the validation dataset based on the micro-F1 score. Since the dataset is highly unbalanced, we use weighted cross-entropy loss for optimization, where the weights are the ratio of validation dataset label distribution to training dataset label distribution, followed by a normalization to ensure that the sum of weights is 1. We use Adam (Kingma and Ba, 2014) optimizer with a learning rate of 0.0005 and batch size of 64. We clip the norm of gradients to 5. We trained our model 6 epochs. The learning rate is annealed by a factor of 0.2 every epoch after epoch 5. We also freeze the embedding for the first two epochs. We use dropout rate of 0.5 in BiLSTM and 0.7 in linear layers. The model is implemented in PyTorch.

4 Result Analysis

In this section we explored different word and sentence representations and compared their performance on the test set. We also conducted sensitivity analysis for our model hyper-parameters. All results are averaged across 5 different seeds. It is worth noting that the settings with the best test scores in the analysis below are not exactly the same as our best system on the leaderboard since our best system is fine-tuned on the validation dataset, which do not guarantee to produce the best test results.

Figure 2: Sensitivity analysis on model hyper-parameters.

Word Representation	micro-F1
Original Glove	0.7250
Pre-trained Glove	0.7279
Fine-tuned Glove	0.7339
ELMo	0.6990
BERT	0.6656

Table 2: Micro-F1 score on the test set using different word representations.

Sentence Representation	micro-F1
None	0.7194
InferSent (GloVe)	0.7259
InferSent (fastText)	0.7277
DeepMoji	0.7299
DeepMoji + InferSent (GloVe)	0.7298
DeepMoji + InferSent (fastText)	0.7344

Table 3: Micro-F1 score on the test set using different sentence representations.

We explored the original GloVe embedding trained on 27B tweet tokens³, pre-trained GloVe embedding⁴, our fine-tuned GloVe embedding, ELMo embedding and BERT embedding. The results are shown in Table 2. Fine-tuned GloVe embedding performs noticeably better than the original GloVe embedding and the pre-trained GloVe embedding. Surprisingly, contextualized embeddings such as ELMo and BERT perform worse than the original GloVe embedding. Possible reasons for their inferior performance are 1) they are fixed during training, which may hinder the overall optimization. 2) they have large embedding size, which can easily cause overfitting.

We explored no sentence embedding, InferSent trained on GloVe, InferSent trained on fastText, and DeepMoji. The results are shown in Table

3. It is clear that sentence representations improved model performance significantly. In particular, DeepMoji achieves the best performance for single sentence representation. InferSent trained on fastText consistently performs better than InferSent trained on GloVe. In addition, concatenating two sentence representations together further improved model performance.

We conducted sensitivity analysis on our model hyper-parameters: hidden size, number of layers in BiLSTM, batch size, learning rate, dropout in BiLSTM and dropout in linear layers. The results are depicted in Figure 2. Our model is relatively robust to hyper-parameters except for the learning rate. When learning rate is around 0.0001 or smaller, the model is unable to be trained effectively.

5 Conclusion

In this paper we presented our model on the task of emotion detection in textual conversations in SemEval-2019. We explored different word and sentence representations in the RCNN model and achieved competitive results. Our result analysis indicate that both pre-trained word and sentence representations help improve the performance of RCNN. However, currently popular contextualized word representations such as ELMo and BERT produced inferior results.

Future improvements can be made on the model architecture. In particular, simply concatenating three utterances into one sentence is not an information-preserving way to incorporate context information. We can design models that can handle a list of utterances and only classify the last utterance to optimize information flow.

³<https://nlp.stanford.edu/projects/glove/>

⁴<https://github.com/cbaziotis/datastories-semeval2017-task4>

References

- Christos Baziotis, Nikos Pelekis, and Christos Douk-
eridis. 2017. Datastories at semeval-2017 task
4: Deep lstm with attention for message-level and
topic-based sentiment analysis. In *SemEval-2017*,
pages 747–754, Vancouver, Canada.
- Ankush Chatterjee, Umang Gupta, Manoj Kumar
Chinnakotla, Radhakrishnan Srikanth, Michel Gal-
ley, and Puneet Agrawal. 2019a. Understanding
emotions in text using deep learning and big data.
Computers in Human Behavior, 93:309–317.
- Ankush Chatterjee, Kedhar Nath Narahari, Meghana
Joshi, and Puneet Agrawal. 2019b. Semeval-2019
task 3: Emocontext: Contextual emotion detection
in text. In *SemEval-2019*, Minneapolis, Minnesota.
- Mathieu Cliche. 2017. Bb.twtr at semeval-2017 task
4: Twitter sentiment analysis with cnns and lstms.
arXiv preprint arXiv:1704.06125.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc
Barraut, and Antoine Bordes. 2017. Supervised
learning of universal sentence representations from
natural language inference data. In *EMNLP*, pages
670–680, Copenhagen, Denmark.
- Laurence Devillers, Ioana Vasilescu, and Lori Lamel.
2002. Annotation and detection of emotion in a
task-oriented human-human dialog corpus. In *pro-
ceedings of ISLE Workshop*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and
Kristina Toutanova. 2018. Bert: Pre-training of deep
bidirectional transformers for language understand-
ing. *arXiv preprint arXiv:1810.04805*.
- Paul Ekman. 1992. An argument for basic emotions.
Cognition & emotion, 6(3-4):169–200.
- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad
Rahwan, and Sune Lehmann. 2017. Using millions
of emoji occurrences to learn any-domain represen-
tations for detecting sentiment, emotion and sar-
casm. In *EMNLP*, pages 1615–1625.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. Twit-
ter sentiment classification using distant supervision.
CS224N Project Report, Stanford, 1(12).
- Scott Gray, Alec Radford, and Diederik P Kingma.
2017. Gpu kernels for block-sparse weights. Tech-
nical report, Technical report, OpenAI.
- Devamanyu Hazarika, Soujanya Poria, Amir Zadeh,
Erik Cambria, Louis-Philippe Morency, and Roger
Zimmermann. 2018. Conversational memory net-
work for emotion recognition in dyadic dialogue
videos. In *NAACL*, volume 1, pages 2122–2132.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997.
Long short-term memory. *Neural computation*,
9(8):1735–1780.
- Jeremy Howard and Sebastian Ruder. 2018. Universal
language model fine-tuning for text classification. In
ACL, volume 1, pages 328–339.
- Rie Johnson and Tong Zhang. 2017. Deep pyramid
convolutional neural networks for text categoriza-
tion. In *ACL*, volume 1, pages 562–570.
- Yoon Kim. 2014. Convolutional neural networks for
sentence classification. In *EMNLP*, pages 1746–
1751.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A
method for stochastic optimization. *arXiv preprint
arXiv:1412.6980*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov,
Richard Zemel, Raquel Urtasun, Antonio Torralba,
and Sanja Fidler. 2015. Skip-thought vectors. In
NIPS, pages 3294–3302.
- Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015.
Recurrent convolutional neural networks for text
classification. In *AAAI*, volume 333, pages 2267–
2273.
- Fei Liu, Trevor Cohn, and Timothy Baldwin. 2018. Re-
current entity networks with delayed memory up-
date for targeted aspect-based sentiment analysis. In
NAACL, volume 2, pages 278–283.
- Navonil Majumder, Soujanya Poria, Devamanyu Haz-
arika, Rada Mihalcea, Alexander Gelbukh, and Erik
Cambria. 2019. Dialoguernn: An attentive rnn for
emotion detection in conversations. In *AAAI*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Cor-
rado, and Jeff Dean. 2013. Distributed representa-
tions of words and phrases and their compositionality.
In *NIPS*, pages 3111–3119.
- Bo Pang, Lillian Lee, et al. 2008. Opinion mining and
sentiment analysis. *Foundations and Trends® in In-
formation Retrieval*, 2(1–2):1–135.
- Jeffrey Pennington, Richard Socher, and Christopher
Manning. 2014. Glove: Global vectors for word
representation. In *EMNLP*, pages 1532–1543.
- Christian S Perone, Roberto Silveira, and Thomas S
Paula. 2018. Evaluation of sentence embeddings
in downstream and linguistic probing tasks. *arXiv
preprint arXiv:1806.06259*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt
Gardner, Christopher Clark, Kenton Lee, and Luke
Zettlemoyer. 2018. Deep contextualized word repre-
sentations. In *NAACL*, volume 1, pages 2227–2237.
- Dominik Scherer, Andreas Müller, and Sven Behnke.
2010. Evaluation of pooling operations in con-
volutional architectures for object recognition. In
ICANN, pages 92–101. Springer.
- Peixiang Zhong, Di Wang, and Chunyan Miao. 2019.
An affect-rich neural conversational model with bi-
ased attention and weighted cross-entropy loss. In
AAAI.