

EmotionX-DLC: Self-Attentive BiLSTM for Detecting Sequential Emotions in Dialogues

Linkai Luo
Hang Seng
Management College
l1k1896@gmail.com

Haiqin Yang
Hang Seng
Management College
hgyang@ieee.org

Francis Y. L. Chin
Hang Seng
Management College
francischin@hsmc.edu.hk

Abstract

In this paper, we propose a self-attentive bidirectional long short-term memory (SA-BiLSTM) network to predict multiple emotions for the EmotionX challenge. The BiLSTM exhibits the power of modeling the word dependencies, and extracting the most relevant features for emotion classification. Building on top of BiLSTM, the self-attentive network can model the contextual dependencies between utterances which are helpful for classifying the ambiguous emotions. We achieve 59.6 and 55.0 unweighted accuracy scores in the *Friends* and the *EmotionPush* test sets, respectively.

1 Introduction

Emotion detection plays a crucial role in developing a smart dialogue system such as a chat conversational bot (Chen et al., 2018). As a typical sub-problem of sentence classification, emotion classification requires not only to understand sentence of a single utterance, but also capture the contextual information from the whole conversations.

The problems of sentence-level classification have been investigated heavily by means of deep neural networks, such as convolutional neural networks (CNN) (Kim, 2014), long short-term memory (LSTM) (Liu et al., 2016), and attention-based CNN (Kim et al., 2018). Additional soft attention layers (Bahdanau et al., 2014) are usually built on top of those networks, such that more attention will be paid to the most relevant words that lead to a better understanding of the sentence. LSTMs (Hochreiter and Schmidhuber, 1997) are also useful to model contextual dependencies. For example, a contextual LSTM model is proposed

to select the next sentence based on the former context (Ghosh et al., 2016), and a bidirectional LSTM (BiLSTM) is adopted to detect multiple emotions (Chen et al., 2018).

In this work, we utilize the self-attentive BiLSTM (SA-BiLSTM) model to predict multiple types of emotions for the given utterances in the dialogues. Our model imitates human’s two-step procedures for classifying an utterance within the context, i.e., sentence understanding and contextual utterances dependence extraction. More specifically, we propose the bidirectional long short-term memory (BiLSTM) with the max-pooling architecture to embed the sentence into a fixed-size vector, as the BiLSTM network is capable of modeling the word dependencies in the sentence while the max-pooling helps to reduce the model size and obtains the most related features for emotion classification. Since data in this challenge is limited and specific words play significant role to classifying the corresponding emotion, we apply the self-attention network (Vaswani et al., 2017) to extract the dependence of all the utterances in the dialogue. Technically, the self-attention model computes the influence of utterance pairs and outputs the sentence embedding of one utterance by a weighted sum over all the utterances in the dialogue. The fully connected layers are then applied on the output sentence embedding to classify the corresponding emotion.

2 Model

Figure 1 presents our designed model architecture. First, the pre-trained 300-dimensional GloVe vectors (Pennington et al., 2014) are adopted to represent each word (token). A sentence (utterance) with m tokens is then represented by

$$S = (w_1, w_2 \dots, w_m), \quad (1)$$

where w_i is d -dimensional word embedding for the i -th tokens in the sentence.

Suppose a dialogue consists of n sentences, the input forms an $n \times m \times d$ tensor, M , see Figure 1. Via the process of sentence embedding (elaborated in Section 2.1), the tensor is converted to a $n \times 2l$ matrix U , where l is the number of the hidden units for each unidirectional LSTM. By applying the self-attentive network, we re-weight the sentence embedding matrix to U' with the same shape as U . Finally, fully connected layers are trained to establish the mapping between input U' and the output emotion labels.

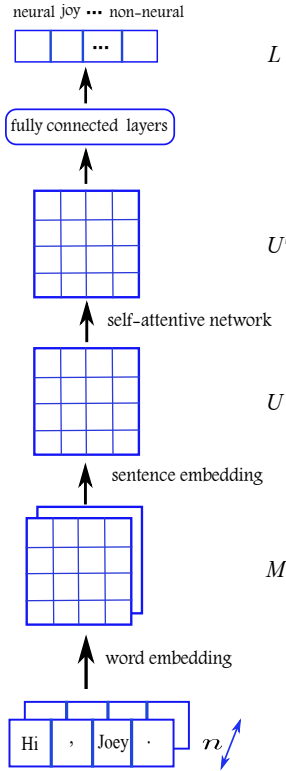


Figure 1: The model architecture illustration. A bunch of n utterances in one dialogue are processed through word embedding, sentence embedding, self-attentive and fully connected layers.

2.1 Sentence Embedding

In this work, we adopt the BiLSTM to learn the sentence embedding because it is the most popular neural network architecture to encode sentences (Conneau et al., 2017; Lin et al., 2017). The forward LSTM and backward LSTM read the sen-




Figure 2: BiLSTM with max-pooling network.

tence S in two opposite directions (see Figure 2):

$$\vec{h}_t = \overrightarrow{LSTM}(w_t, \vec{h}_{t-1}) \quad (2)$$

$$\overleftarrow{h}_t = \overleftarrow{LSTM}(w_t, \overleftarrow{h}_{t+1}) \quad (3)$$

The vectors \vec{h}_t and \overleftarrow{h}_t are concatenated to a hidden state h_t . Max-pooling (Collobert and Weston, 2008) is then conducted along all the words of a sentence to output the final sentence representation, u .

2.2 The Self-attentive Network

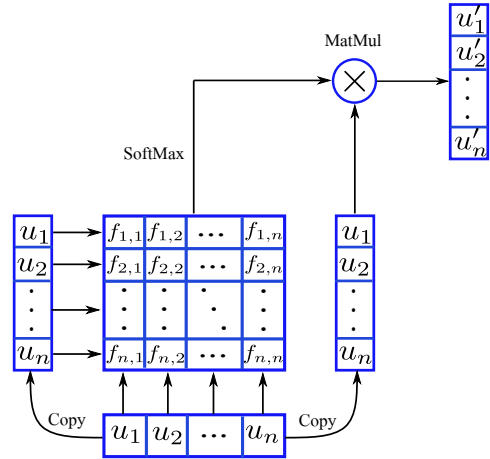


Figure 3: The self-attentive network. f_{ij} denotes $f(u_i, u_j)$ in Eq. 4.

The self-attention model, called Transformer (Vaswani et al., 2017), is an effective non-recurrent architecture for machine translation. We adopt it to capture the utterances

dependence. Figure 3 shows the model to build the dot-product attention of utterances, where the attention matrix is calculated by:

$$f(u_i, u_j) = \begin{cases} \frac{u_i u_j^T}{\sqrt{d_k}}, & \text{if } i, j \leq n; \\ -\infty & \text{otherwise.} \end{cases} \quad (4)$$

where u_i is the i -th sentence embedding in the dialogue, and d_k (i.e. $2l$) is the model dimension. An attention mask is applied to waive the inner attention between sentence embeddings and paddings.

The i -th sentence embedding is finally weighted by summing over all the sentence embeddings to enhance the effect:

$$u'_i = \sum_j \text{softmax}(F_i) u_j, \quad (5)$$

where F_i is n -dimensional vector whose j -th element is $f(u_i, u_j)$.

2.3 Output and Loss

Finally, we apply fully connected layers to produce the corresponding emotions. In the training, the weighted cross-entropy is adopted as the loss function. Since the challenge only focuses on classifying four types of emotions, rather than all eight types, we set the weights to zero for the unconsidered emotions.

3 Experiments and Results

3.1 Data Preprocessing

The EmotionX dataset consists of the *Friends* TV scripts and the *EmotionPush* chat logs on Facebook Messenger in eight types, i.e., Neutral, Joy, Sadness, Anger, Fear, Surprise, Disgust and Non-neutral. For the train set, there are 720 dialogues for *Friends* and *EmotionPush*, respectively, which yields a total of 1,440 dialogues, 21,294 sentences, and 9,885 unique words. In the challenge, we test the following candidate labels, Neutral, Joy, Sadness, and Anger. We also conduct the following steps to clean the data:

- Unicode symbols, except emojis (the direct expressions of human emotions), are removed. Person names, locations, numbers and websites are replaced with special tokens.
- The Emoji symbols are converted to the corresponding meanings.

- Duplicated punctuation and symbols. Tokens with duplicated punctuation or alphabets, such as “oooooh”, often imply non-neural emotions. We reconstruct the tokens to be oh <duplicate> to avoid informal words. The same rule also applies to similar tokens. For example, “oh!!!!!!” is replaced by oh ! <duplicate>.

- Word tokenization. We use NLTK’s Twitter-Tokenizer (Bird, 2006) to split the sentences into tokens. All tokens are set lowercase.

3.2 Experimental Setup

We conduct two experiments with different model variants: BiLSTM and SA-BiLSTM, to validate whether our proposed model can learn the contextual information. The network settings for each model are summarized as follows:

- BiLSTM: BiLSTM + max-pooling + fully connected layers.
- SA-BiLSTM: BiLSTM + max-pooling + self-attentive network + fully connected layers.

The word embedding is 300-dimensional from the Glove. Pack padded sequence and pad packed sequence are implemented to deal with varying sequence lengths. For SA-BiLSTM, we limit the utterance number to 25 for each dialogue. Due to the limit of training data, LSTM is set to one layer with only 256 hidden units. The fully connected layers consist of two middle layers with the same size of 128.

The mini-batch size for training BiLSTM is set to 16. Unlike BiLSTM, we feed one dialogue to SA-BiLSTM for every training step. Adam (Kingma and Ba, 2014) is the adopted optimizer with initial learning rate 0.0002 and decay factor 0.99 for every epoch. Dropout probability is set to 0.3 for BiLSTM and self-attention layers. We train BiLSTM for 10 epochs and SA-BiLSTM for 20 epochs to gain the best accuracy in the validation sets.

3.3 BiLSTM Versus SA-BiLSTM

Table 1 reports the model performance in the validation sets, which consist of 80 dialogues for *Friends* and *EmotionPush*, respectively. We evaluate two criteria, the weighted accuracy (WA) and the unweighted accuracy (UWA) (Chen et al.,

Model	Dataset	WA	UWA	Neutral	Joy	Sadness	Anger
BiLSTM	<i>Friends</i>	79.4	60.4	92.5	78.9	29.0	41.2
	<i>EmotionPush</i>	83.9	61.8	92.6	73.1	41.0	40.4
SA-BiLSTM	<i>Friends</i>	78.8	62.8	90.6	73.2	40.3	47.1
	<i>EmotionPush</i>	83.4	63.5	91.9	69.6	47.0	45.7

Table 1: Experimental results of *Friends* and *EmotionPush* in the validation sets.

Model	Dataset	UWA	Neutral	Joy	Sadness	Anger
SA-BiLSTM	<i>Friends</i>	59.6	90.1	68.8	30.6	49.1
	<i>EmotionPush</i>	55.0	94.2	70.5	31.0	24.3
	<i>Average</i>	57.3	92.1	69.6	30.8	36.7

Table 2: Experimental results of *Friends* and *EmotionPush* in the test sets.

2018). The predicted accuracy for each class is also given in the table.

Interestingly, the simpler model BiLSTM achieves higher WA, with up to 0.6% and 0.5% improvement in *Friends* and *EmotionPush*, respectively. On the other hand, SA-BiLSTM outperforms BiLSTM in terms of UWA, with up to 1.4% and 1.7% improvement. Note that BiLSTM tends to predict the emotions Neutral and Joy far more accurately than the other two emotions because most utterances are labeled as these two emotions, i.e., 45.03% as neutral and 11.79% as joy in *Friends* while 66.85% as neutral and 14.25% as joy in *EmotionPush*. Overall, SA-BiLSTM provides a more balanced prediction for each type of emotion than BiLSTM. Especially in predicting the emotions of Sadness and Anger, SA-BiLSTM gains better predictive accuracy, up to 11.3% & 6.0% on the Sadness emotion and 5.9% & 5.3% on the Anger emotion improvements in *Friends* and *EmotionPush*, respectively.

3.4 Results of Test Set

We submit the results produced by SA-BiLSTM and obtain the evaluation scores provided by the challenge organizer. Table 2 lists the experimental results evaluated in the test set, which consists of 400 dialogues, 200 dialogues for *Friends* and *EmotionPush*, respectively.

The results indicate that our model shows a strong bias towards predicting the Neutral emotion and the Joy emotion in both datasets compared to the Sadness and the Anger emotions. Especially, our model achieves an extremely poor prediction on the Anger emotion in *EmotionPush*. Moreover, the UWA in *EmotionPush* is smaller than that in

Friends, which is different from our prediction results in the validation set. We conjecture that the distribution of the validation set and the test set may be slightly different. To obtain a robust solution, we may train multiple models using different random seeds and ensemble the model averaged on the checkpoints.

We notice that the *Speaker* information is also important for emotion classification. Table 3 shows two consecutive utterances made by the same speaker from *EmotionPush*, where the first utterance seems literally less emotional than the second one. Nevertheless, the two utterances should carry the same emotion, i.e., Anger, because they are made by the same speaker consecutively. On the contrary, our model gives a false prediction (i.e., Neutral) for the second utterance because it probably treats the two utterances separately. We believe that our model shall gain some improvements by adding speaker information into it.

Speaker ID	Utterance
1051336806	<i>but /you/ bug /me/</i>
1051336806	<i>and you hundred percent told peopel stfu.</i>

Table 3: Consecutive utterances made by the same speaker shall carry the same emotion.

4 Conclusion

In this work, we propose SA-BiLSTM to predict multiple emotions for given utterances in the dialogues. The proposed network is a self-attentive network built on top of BiLSTM. Our results evaluated on the validation set show that BiLSTM has

better WA performance, while SA-BiLSTM is advantageous to BiLSTM in terms of UWA. According to the test results, SA-BiLSTM yields higher UWA scores for detecting the Neural emotion and the Joy emotions than the Sadness and the Anger ones. The bias may be caused by uneven training data distributions. We hope to improve our model by either incorporating more related data or retrieving more linguistic information.

Acknowledgment

This work was partially supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. UGC/IDS14/16).

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *CoRR*, abs/1409.0473.
- Steven Bird. 2006. [NLTK: the natural language toolkit](#). In *ACL 2006, 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, Sydney, Australia, 17-21 July 2006*.
- Sheng-Yeh Chen, Chao-Chun Hsu, Chuan-Chun Kuo, Ting-Hao Huang, and Lun-Wei Ku. 2018. [Emotionlines: An emotion corpus of multi-party conversations](#). *CoRR*, abs/1802.08379.
- Ronan Collobert and Jason Weston. 2008. [A unified architecture for natural language processing: deep neural networks with multitask learning](#). In *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5-9, 2008*, pages 160–167.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 670–680.
- Shalini Ghosh, Oriol Vinyals, Brian Strope, Scott Roy, Tom Dean, and Larry P. Heck. 2016. [Contextual LSTM \(CLSTM\) models for large scale NLP tasks](#). *CoRR*, abs/1602.06291.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Yanghoon Kim, Hwanhee Lee, and Kyomin Jung. 2018. [Attnconvnet at semeval-2018 task 1: Attention-based convolutional neural networks for multi-label emotion classification](#). *CoRR*, abs/1804.00831.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1746–1751.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Zhouhan Lin, Minwei Feng, Cícero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. [A structured self-attentive sentence embedding](#). *CoRR*, abs/1703.03130.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. [Recurrent neural network for text classification with multi-task learning](#). In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 2873–2879.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6000–6010.