

A Convolutional Attention Model for Text Classification

Jiachen Du^{1,2}, Lin Gui¹, Ruifeng Xu^{1,3*}, Yulan He⁴

¹ Laboratory of Network Oriented Intelligent Computation,
Shenzhen Graduate School, Harbin Institute of Technology, Shenzhen, China

² Department of Computing, the Hong Kong Polytechnic University, Hong Kong

³ Guangdong Provincial Engineering Technology Research Center for Data Science

⁴ School of Engineering and Applied Science, Aston University, United Kingdom
dujiachen@stmail.hitsz.edu.cn, xuruifeng@hit.edu.cn, y.he9@aston.ac.uk,
guilin.nlp@gmail.com

Abstract. Neural network models with attention mechanism have shown their efficiencies on various tasks. However, there is little research work on attention mechanism for text classification and existing attention model for text classification lacks of cognitive intuition and mathematical explanation. In this paper, we propose a new architecture of neural network based on the attention model for text classification. In particular, we show that the convolutional neural network (CNN) is a reasonable model for extracting attentions from text sequences in mathematics. We then propose a novel attention model base on CNN and introduce a new network architecture which combines recurrent neural network with our CNN-based attention model. Experimental results on five datasets show that our proposed models can accurately capture the salient parts of sentences to improve the performance of text classification.

1 Introduction

In recent years, there is no doubt that deep learning has ushered in amazing technological advances on natural language processing(NLP) researches. Much of the work with deep learning involves learning word vector representations through natural language models [15] and composition over word vectors for various tasks like text classification [4], machine translation [9], document summarization [2] and so on.

In this paper, we focus on the text classification problem. Traditional approaches to text classification firstly represent text sequences with sparse features, such as n-grams [3], topic-based representation [26] and kernel methods [8]. Recently deep-learning models have shown their big success in text classification, such as convolutional neural networks [12] and recurrent neural networks based on long short-term memory [7].

Applying convolutional neural network (CNN) to NLP including text classification has drawn many interests in recent years. It has been shown that CNNs can be directly applied to embeddings of words [11]. Unlike word level models, [27] proposed a character-level CNN for text classification which achieved competitive results. Although CNN has been proven efficient on text classification, it usually ignores important

* Corresponding author: Ruifeng Xu

long-distance sequential information which greatly impacts the classification performance, especially in sentences which have negation and semantic transition. Recurrent neural network (RNN) is another important model in NLP. [22] used gated recurrent neural network to model documents and applied this model to sentiment classification. [21] explored the structure of a sentence and used a tree-structured recurrent neural network with long-short term memory (LSTM) for text classification. The advantage of RNN is its ability to better capture the contextual information, especially the semantics of long texts. However RNN model cannot pay attention to the salient parts of text. This limitation reduces the effectiveness of RNN when applied to text classification.

Recently, based on the aforementioned architectures, a new direction of neural networks has emerged. It learns to focus "attention" to specific parts of text as the simulation of human's attention while reading. However, the researches on neural network with attention mechanism only show promising results on a sequence-to-sequence (seq2seq) tasks in NLP, including machine translation [1], caption generation [24] and text summarization [18]. It's not appropriate to use the same alignment mechanism for classification. For example, [25] applies the attention model used in seq2seq tasks to document-level classification. The attention is modeled by a one hidden layer perceptron with RNN hidden unit as input. Compare with non-attention methods, the improvement is limited in the reported result. The reason is that, in the seq2seq problem, each word has a corresponding label, such as the word in another language (for machine translation tasks), or a Part-Of-Speech label (for POS tagging tasks). In text classification, there are no target labels for each word to indicate the word is category-relevant or not. So, there is no evidence showing why the perceptron with one hidden layer attention is efficient for text classification.

Motivated by the cognitive and neuroscience research, we propose a novel attention model. The basic idea is to first use the convolution operation to capture attention signals, each of which stands for the local information of a word in its context; then use RNN to model text with attention signals. A word with higher attention weight, which carries more valuable information, will be more important in text modeling. Our main contributions in this paper is three-fold:

- This is the first time a convolutional neural network model is presented to stimulate human's reading attention based on cognitive and neuroscience research, and a detailed mathematical deduction of this model is given in this work.
- Moreover, a novel attention extraction method based on this model has been proposed that can be used in several tasks.
- Finally, we propose a new architecture based on convolutional attention extraction model, and this neural network shows competitive results compared with state-of-the-art models in text classification.

The rest of our paper is structured as follows, Section 2 explores a novel attention model based on convolutional neural networks and gives detailed reduction of this model. Section 3 gives a description of the Convolutional Recurrent Attention Network (CRAN) for text classification. Section 4 presents some experiments to justify the effectiveness of CRAN on text classification. Section 5 concludes the paper and outlines the future work.

2 A CNN-based Attention Model

It is commonly known that we can pay attention to only small amount of information presented in visual scenes and only concentrate on the information related to a specific task at hand. Cognitive and neuroscience researches have confirmed this hypothesis, and a lot of experiments have shown that humans depict in brains a visual representation or "search template" of certain task and try to only pay attention to the information which can match the "task-oriented template" [5]. psycholinguistics has proven that template-matching process also helps us emphasize the important content while our brain is processing texts. Although this mechanism of attention has been thoroughly studied in neuroscience and psychology, there is few research on how to introduce it to computational linguistic and natural language processing. Motivated by that, we propose a novel model introducing the aforementioned attention mechanism to natural language processing , especially to text classification.

Based on in-depth investigation, we found that Convolutional Neural Network (CNN) is a natural model to stimulate human being's reading attention mechanism , since the convolution operation as the core component of CNN is similar to the process of template matching. For textual data, CNN always applies one-dimensional convolution to the concatenate of vector representation of each words. Our first goal is to show that one-dimension convolution of CNN is precisely the process of calculating the similarity between snippets of text and the "attention searching templates". In neural-network based models, a text sequence of length T (padded where necessary) is often represented as

$$x_{0:T-1} = x_0 \oplus x_1 \oplus \dots \oplus x_{T-1} \quad (1)$$

where $x_t \in \mathcal{R}^d$ ($t = \{0, 1, \dots, T-1\}$) corresponds to the d -dimensional vector representation of the t -th word in the text sequence, and \oplus is the vector concatenation operator. Convolution operation applies a filter $w = w_0 \oplus w_1 \oplus \dots \oplus w_{l-1}$ to a window of l words in the original sequence to get a convolutional similarity c_t , where each column of the filter $w_t \in \mathcal{R}^d$ ($t = \{0, 1, \dots, l-1\}$) is a vector of the same length of word embeddings. The convolution operation apply the following transformation to each continuous subsequence of length l in $[x_0, x_1, \dots, x_{T-1}]$. Suppose the current subsequence is $x_{t:t+l-1} = x_t \oplus x_{t+1} \oplus \dots \oplus x_{t+l-1}$:

$$\begin{aligned} c_t &= f(\langle w, x_{t:t+l-1} \rangle + b) \\ &= f\left(\|w\| \times \|x_{t:t+l-1}\| \times \frac{\langle w, x_{t:t+l-1} \rangle}{\|w\| \times \|x_{t:t+l-1}\|} + b\right) \end{aligned} \quad (2)$$

In equation 2, $\langle \cdot, \cdot \rangle$ is dot product of two vectors as $\langle a, b \rangle = a^T b$, $\|\cdot\|$ is the \mathcal{L}^2 norm of vectors. f is the non-linear function such as hyperbolic tangent, sigmoid and rectified linear function. Notice that w and $x_{t:t+l-1}$ are $l \times d$ -dimensional vector, assuming that each dimension has its own distribution. According to Chebyshev Law, there exists M , for any w and any $x_{t:t+l-1}$, if $l \times d$ is large enough (usually larger than 25), it is true that for any $\varepsilon > 0$, $P(|M - \|w\| \times \|x_{t:t+l-1}\|| < \varepsilon) = 1$. since $l \times d$, namely the shape of convolution filters is always larger than 25, we can replace

$\|w\| \times \|x_{t:t+l-1}\|$ by M in this equation, and define a function $F(x) = f(Mx)$ to replace the original function f , let $b' = b/M$, we will obtain:

$$c_t = F(\cos(w, x_{t:t+l-1}) + b') \quad (3)$$

In equation 3 we notice that F is only a compression function that satisfies $F'(x) > 0$. Then we can consider the convolutional filter w be the "search template" mentioned before in human's attention while reading. And c_t can be treated as the cosine similarity between the search template and the part of text which is processed currently. b' in equation 3 is the threshold of this similarity. When the similarity is greater than b' , the textual part being processed is considered task-relevant; otherwise it is task-irrelevant.

The output of each convolutional filter is an attention signal from the original text. In order to reduce the noise in the attention signals, we choose multiple convolutional filters applied to the vector representation of original text and average the results of these filters to obtain a smooth attention signal. Suppose the number of convolutional filters is m , these filter is denoted by $[w^1, w^2, \dots, w^m]$, and the corresponding attention similarity is $[c^1, c^2, \dots, c^m]$. After averaging the attention similarities along the filter-axis, we will obtain the smooth attention signal $c \in \mathcal{R}^T$ of which each element represents the importance of the corresponding word.

$$c = \sum_{i=1}^m c^i \quad (4)$$

The whole CNN-based attention model is shown in Figure 1.

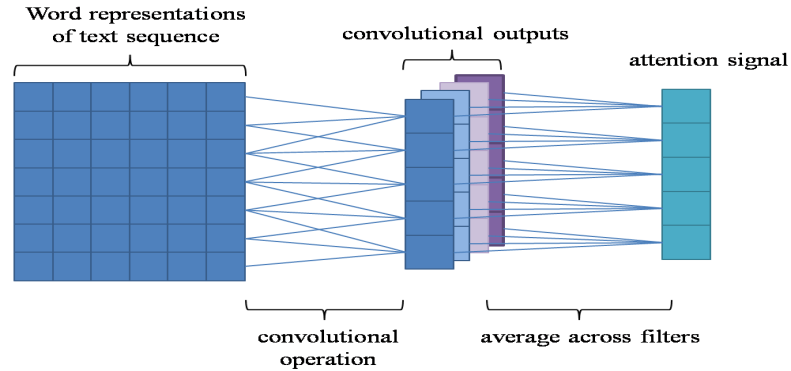


Fig. 1. CNN-based model of attention extraction.

3 Convolutional-Recurrent Attention Neural Networks

Based on the attention model we described in Section 2, we propose a model named Convolutional-Recurrent Attention Network (CRAN) that combines recurrent neural

network (RNN) and the convolutional attention model in Section 2. The reason why we use RNN as a part of our model rather than directly using the traditional CNN architecture is that traditional CNN uses a pooling layer over the whole sentence which results in a single vector being extracted as the representation of the sequence. It makes CNN difficult to capture the long-distance dependencies in sequences like negation and transition. However RNN with long short-term memory (LSTM) is designed for handling the long-distance dependency problem. We speculate that combining RNN with our proposed CNN-based attention model will give better performance compared with using a pure CNN model, and the experimental results which will be presented in Section 4 confirm our hypothesis.

The overall architecture of the Convolutional Attention Neural Network (CRAN) is shown in Figure 2. It consists of two main parts: a recurrent neural network (RNN) as the text encoder and a convolutional neural network (CNN) as the attention extractor. We describe the details of these two parts in the following subsections.

3.1 RNN-Based Sequence Encoder

An RNN [6] is a kind of neural network that processes sequences of arbitrary length by recursively applying a function to its hidden state vector $h_t \in \mathcal{R}^d$ of each element in the input sequences. The hidden state vector at time-step depends on the input symbol x_t and the hidden state vector at last time-step h_{t-1} is:

$$h_t = \begin{cases} 0 & t = 0 \\ g(g_{t-1}, x_t) & \text{otherwise} \end{cases} \quad (5)$$

A fundamental problem in traditional RNN is that gradients propagated over many steps tend to either vanish or explode. It makes RNN difficult to learn long-dependency correlations in a sequence. Long short-term memory network (LSTM) was proposed by [7] to alleviate this problem. LSTM has three gates: an input gate i_t , a forget gate f_t , an output gate o_t and a memory cell c_t . They are all vectors in \mathcal{R}^d . The LSTM transition equations are:

$$\begin{aligned} i_t &= \sigma(W_i x_t + U_i h_{t-1} + V_i c_{t-1}), \\ f_t &= \sigma(W_f x_t + U_f h_{t-1} + V_f c_{t-1}), \\ o_t &= \sigma(W_o x_t + U_o h_{t-1} + V_o c_{t-1}), \\ \tilde{c}_t &= \tanh(W_c x_t + U_c h_{t-1}), \\ c_t &= f_t \odot c_{t-1} + i_t \odot \tilde{c}_t, \\ h_t &= o_t \odot \tanh(c_t) \end{aligned} \quad (6)$$

where x_t is the input at the current time step, σ is the sigmoid function and \odot is the elementwise multiplication operation. In our model, we use the hidden-state vector of each time step as the representation of corresponding word in the sentence.

3.2 CNN-Based Attention Extraction

As discussed in Section 2, we use a CNN-based network to model the attention signal in sentences. Suppose the input text sequence is $[x_0, x_1, \dots, x_{T-1}]$, where $x_t \in \mathcal{R}^d$ ($t =$

$0, 1, \dots, T-1$), m convolutional filters of length l are denoted as $[w^1, w^2, \dots, w^m]$, the corresponding convolution results are $[c^1, c^2, \dots, c^m]$, where $c^i \in \mathcal{R}^T$ ($t = 1, 2, \dots, m$) represents the attention distributed on the sequence of length T . We average the m filters to get the stable attention signal c which is a vector of length T as described in Section 2.

3.3 Text Classification

We use the product of attention signal c_t and the corresponding hidden state vector of RNN h_t to represent the word t in a sequence with attention signal. The representation of the whole sequence can be obtained by averaging the word representations:

$$s = \frac{1}{T} \sum_{t=0}^{T-1} c_t h_t \quad (7)$$

where $s \in \mathcal{R}^d$ is the vector representation of the text sequence and it can be used as features for text classification:

$$p = \text{softmax}(W_c s + b_c) \quad (8)$$

where p is the predicted label for text sequence s , W_c and b_c are parameters of the classification layer.

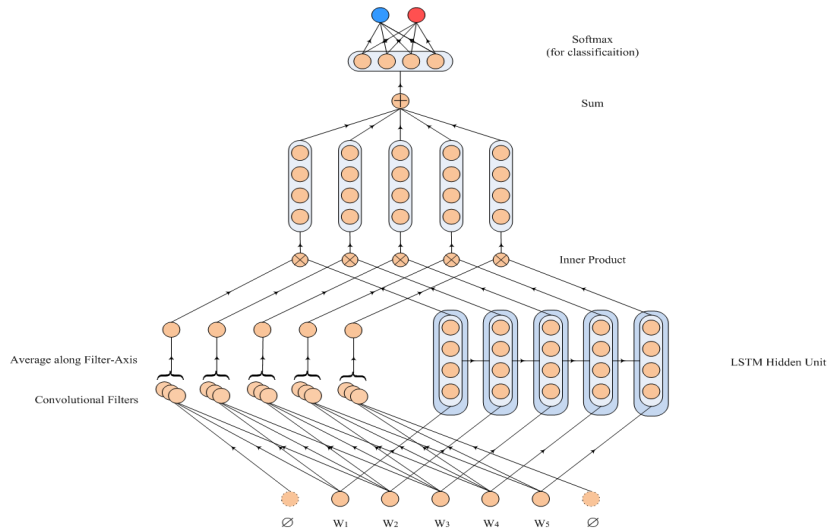


Fig. 2. Architecture of convolutional attention network

4 Experiments

In this section, we investigate the empirical performance of our proposed CRAN on various datasets and compare it with state-of-the-art models for text classification.

4.1 Datasets

We choose five different text classification datasets and test our model on various benchmarks. We briefly summarize these datasets as follows:

- **MR**: Movie reviews with one sentence per review. Classification involves binary categories of reviews (positive and negative) ¹ [17].
- **SST-1**: An extension of MR but with train/dev/test splits provided and fine-grained labels (negative, somewhat negative, neutral, somewhat positive, positive) ² [20].
- **SST-2**: Same as SST-1 but with neutral reviews removed and only containing binary labels (positive and negative).
- **Subj**: Subjectivity dataset where task is to classify a sentence as objective or subjective [16].
- **IMDB**: A document-level text classification dataset containing 100,000 movie reviews with binary labels ³ [14].

Dataset	Class	Average Length	Vocabulary Size	Text Size	Test Size
MR	2	20	18K	10662	CV
SST-1	5	19	18K	11855	2210
SST-2	2	18	15K	9613	1821
Subj	2	21	21K	10000	CV
IMDB	2	294	392K	50000	25000

Table 1. Summary statistics of the five datasets.

The first four datasets are for sentence-level classification and the last one is for document-level classification. We conduct experiments on the standard test sets on SST-1, SST-2 and IMDB. For datasets without standard train/test split, we use 10-fold cross validation instead. The summary statistics of the datasets are listed in Table 1.

4.2 Model Training and Hyper-parameters

We train our proposed convolutional attention model with two different initialization strategies:

- **CRAN_rand** : The convolutional attention layer is randomly initialized.

¹ <http://www.cs.cornell.edu/people/pabo/movie-review-data/>

² <http://ai.stanford.edu/sentiment/>

³ <http://ai.stanford.edu/amaas/data/sentiment/>

- **CRAN_pretrain** : The convolutional attention layer is pre-trained by a standard standard convolutional neural network which is connected to a fully-connected layer to the final labels as a classifier.

The model can be trained in an end-to-end way by back-propagation, where the objective function is cross-entropy of error loss. Training is done through gradient descent with the Adadelta update rule. In all of these experiments, the word embeddings are initialized with the publicly available word2vec vectors that were trained on 100 billion words from Google News (Mikolov, Sutskever, et al., 2013) . Other parameters are set as follows. The number of hidden units of LSTM and convolution filters in CNN is 100 , the length of convolution filter is 3, dropout rate is 0.5 and mini-batch size is 16. These hyper-parameters are chosen via a grid search on the SST-2 development set.

4.3 Baselines

To illustrate the performance boost of our proposed attention model, we compare our model with some baseline methods. Since we use RNN as component of our model, we implement an RNN with LSTM memory unit as a baseline. We also compare our model with LSTM with the attention model proposed by [25].

- **LSTM**: LSTM for text classification.
- **LSTM + RNN attention** attention: The attention-based LSTM model proposed by [25]. Since most datasets used in our experiments are for sentence-level classification, we implement a flatten variant of this model without aggregating the attention signals from sentences to form the document-level attention signal.

We also compare our model with the following state-of-the-art models:

- **NBOW**: The NBOW sums the word vectors within the sentence and applies a softmax classifier.
- **MV-RNN**: Matrix-Vector Recursive Neural Network with parse trees [19].
- **F-Dropout**: Fast Dropout from [23].
- **DCNN**: Convolutional Neural Network with dynamic k-max pooling [10].
- **CNN**: Convolutional neural networks with max pooling [11].
- **Tree-LSTM**: Tree-Structured Long Short-Term Memory Networks [21].
- **Multi-Task**: Shared-layer multi-task learning model trained on four different datasets [13].

4.4 Results and Analysis

The experimental results on all datasets are shown in Table 2. Firstly we notice that LSTM performs the worst among all the models since the sequential features extracted by RNN models is not suitable for text classification. But we find that our proposed CRAN can combine the merits of CNN and RNN to improve the performance. Results show that CRAN improves the performance significantly compared with the LSTM by 3.0% on average. For the two variations of CRAN, we can see that CRAN_rand model performs slightly better than CRAN_pretrain model except on the IMDB dataset. This

Model	MR	SST-1	SST-2	Subj	IMDB
CRAN_rand	82.8	50.0	87.7	94.1	92.0
CRAN_pretrain	82.0	48.1	86.9	94.0	92.1
LSTM	80.1	46.2	85.2	91.2	88.5
LSTM + RNN attention	82.0	48.0	86.1	93.2	90.6
<hr/>					
NBOW	77.1	42.1	79.0	90.8	80.7
MV-RNN(Socher et al., 2012)	79.0	44.4	82.9	-	-
F-Dropout (Wang et al. 2013)	79.1	-	-	93.6	91.1
DCNN(Kalchbrenner et al., 2014)	-	48.5	86.8	-	-
CNN (Kim, 2014)	81.5	48.0	88.1	93.4	-
Tree-LSTM (Tai et al. 2015)*	-	50.6	86.9	-	-
Multi-Task (Liu et al. 2016)*	-	49.6	87.9	94.1	91.3

Table 2. Results of our proposed CRAN against baselines. Results marked * are models that need external tools or resources.

is because CRAN_pretrain uses a different architecture to pre-train the convolutional attention layer. Both CRAN variants perform better than the LSTM with RNN attention. This shows that our proposed attention modelling method can capture more accurate attention information from text.

Compared with the state-of-the-art models in text classification, CRAN gives the best performance on three out of five datasets. Although CRAN uses a similar CNN network as its attention extractor, CRAN improves upon CNN by 0.9% on average. This is because combining RNN helps our model capture the long-distance semantic dependencies in sentences which cannot be dealt with by traditional CNNs. Tree-LSTM outperforms our model on SST-1 by 0.6%. However it needs an external parser to derive the tree-structure of each sentence, and the performance listed in Table 2 is reported on the exact parsing results of sentences labelled by annotators. It is worth noting that our models are comparable with the recurrent neural network with multi-task learning proposed by (Liu et al., 2016). This model is an extremely strong baseline which was trained jointly on four datasets (SST-1, SST-2, Subj and IMDB).

4.5 Case Study

In order to validate that our model is able to select salient parts in a text sequence, we visualize the attention layers in Figure 3 for several sentences from the MR dataset in which our model predicted the class labels correctly. Each line in the figure is a heatmap of the attention signals extracted by our model. The attention signal is normalized in the range of [-1,1], and the actual label of each sentence is shown to the left of each heatmap. The red color in the heatmap means a high attention for the positive label and the blue color means a high attention for the negative label, while the white color means neutral.

Figure 3 shows that our model can not only select words carrying strong sentiments like *repelling*, *annoyance*, *dull*, *et.al.*, but also deal with transitions of sentiments in sentences. For example, in the second sentence, our model assigns a high positive attention to the first half of the sentence before the word *but*, and it also finds the second half

of this sentence is highly related to the negative sentiment. Also, the negative attention value is greater than the positive attention value in this sentence, which makes the final predicted label to be negative.

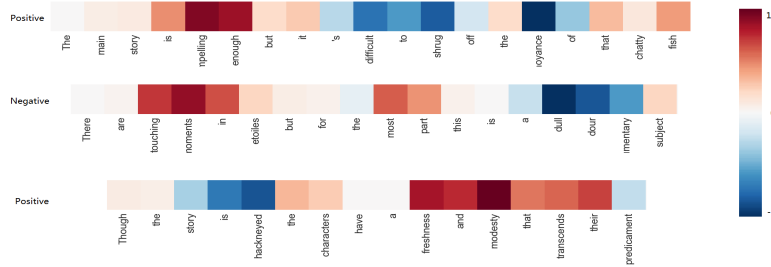


Fig. 3. Visualization of attention signals in sample sentences in the MR dataset.

We also analyze the sentences from which our model failed to extract attention signals properly. We show an attention visualization in Figure 4 which is extracted from a sentence in the MR dataset:

Imagine the James Woods character from video drome making a home movie of audrey rose and showing it to the kid from the sixth sense and youve imagined the ring.

The label of this sentence is negative, but our model predicts its labels positive. We can see that this is a complex sentence with some metaphors. To understand the real meaning of this sentence, readers must have the background knowledge of the movies mentioned in this sentence. For people who are unfamiliar with the movies, James Woods and Sixth Sense, they will not know these are all horror movies and hence would have a difficulty in understanding the metaphoric meaning of these terms.

We can observe that, for this sentence, the attention signals extracted by our model is somewhat randomized and do not reflect the actual sentiment expressed in the sentence. It seems that our model can only extract the attention signals from the literal meanings of words. Nevertheless, understanding complicated linguistic phenomena such as analogy, metaphor and irony is a huge challenge commonly faced by many NLP tasks, not just in text classification.

5 Conclusion

In this paper, we have shown that the convolutional neural network is a reasonable model for extracting attentions from text sequences. Based on this finding, we have proposed a novel attention extraction model based on the convolutional neural network. Utilizing this CNN-based attention model, we have introduced a new neural network architecture which combines RNN with our CNN-based attention model. We have conducted extensive experiments on five datasets. The experimental results show that (1)

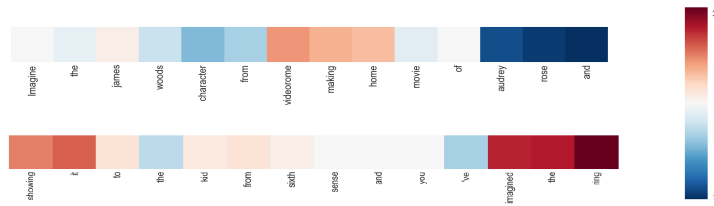


Fig. 4. Visualization of wrongly classified sentence in the MR dataset.

our model is capable of extracting salient parts from sentences; (2) our model can combine the merits of CNN and RNN to improve the sentence classification performance. Finally, the visualization of some attentions extracted by our model shows the impressive capability of our model to process the sentiment transitions in sentences. We have also presented an attentional visualization result of a sentence whose class label was wrongly predicted by our model. This shows that our model has a difficulty in dealing with more subtle meanings embedded in sentences, a huge challenge commonly faced by many NLP tasks. In future works, we will mainly focus on extending our CNN-based attention model to other tasks like text generation and sequence to sequence learning.

Acknowledgements

This work was supported by the National Natural Science Foundation of China 61370165, U1636103, 61632011, 61528302, Shenzhen Foundational Research Funding JCYJ20150625142543470, Guangdong Provincial Engineering Technology Research Center for Data Science 2016KF09 and grant from the Hong Kong Polytechnic University (G-YBJP).

References

1. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014)
2. Cao, Z., Wei, F., Dong, L., Li, S., Zhou, M.: Ranking with recursive neural networks and its application to multi-document summarization. In: AAAI. pp. 2153–2159 (2015)
3. Cavnar, W.B., Trenkle, J.M.: N-gram-based text categorization. Ann Arbor MI 48113(2), 161–175 (1994)
4. Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., Kuksa, P.: Natural language processing (almost) from scratch. Journal of Machine Learning Research 12(Aug), 2493–2537 (2011)
5. Duncan, J., Humphreys, G.W.: Visual search and stimulus similarity. Psychological review 96(3), 433 (1989)
6. Elman, J.L.: Finding structure in time. Cognitive science 14(2), 179–211 (1990)
7. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation 9(8), 1735–1780 (1997)
8. Joachims, T.: Text categorization with support vector machines: Learning with many relevant features. In: European conference on machine learning. pp. 137–142. Springer (1998)

9. Kalchbrenner, N., Blunsom, P.: Recurrent continuous translation models. In: EMNLP. vol. 3, p. 413 (2013)
10. Kalchbrenner, N., Grefenstette, E., Blunsom, P.: A convolutional neural network for modelling sentences. arXiv preprint arXiv:1404.2188 (2014)
11. Kim, Y.: Convolutional neural networks for sentence classification. arXiv preprint arXiv:1408.5882 (2014)
12. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11), 2278–2324 (1998)
13. Liu, P., Qiu, X., Huang, X.: Recurrent neural network for text classification with multi-task learning. arXiv preprint arXiv:1605.05101 (2016)
14. Maas, A.L., Daly, R.E., Pham, P.T., Huang, D., Ng, A.Y., Potts, C.: Learning word vectors for sentiment analysis. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*. pp. 142–150. Association for Computational Linguistics (2011)
15. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in neural information processing systems*. pp. 3111–3119 (2013)
16. Pang, B., Lee, L.: A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In: *Proceedings of the 42nd annual meeting on Association for Computational Linguistics*. p. 271. Association for Computational Linguistics (2004)
17. Pang, B., Lee, L.: Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In: *Proceedings of the 43rd annual meeting on association for computational linguistics*. pp. 115–124. Association for Computational Linguistics (2005)
18. Rush, A.M., Chopra, S., Weston, J.: A neural attention model for abstractive sentence summarization. arXiv preprint arXiv:1509.00685 (2015)
19. Socher, R., Huval, B., Manning, C.D., Ng, A.Y.: Semantic compositionality through recursive matrix-vector spaces. In: *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. pp. 1201–1211. Association for Computational Linguistics (2012)
20. Socher, R., Perelygin, A., Wu, J.Y., Chuang, J., Manning, C.D., Ng, A.Y., Potts, C.: Recursive deep models for semantic compositionality over a sentiment treebank. In: *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*. vol. 1631, p. 1642. Citeseer (2013)
21. Tai, K.S., Socher, R., Manning, C.D.: Improved semantic representations from tree-structured long short-term memory networks. arXiv preprint arXiv:1503.00075 (2015)
22. Tang, D., Qin, B., Liu, T.: Document modeling with gated recurrent neural network for sentiment classification. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing* pp. 1422–1432 (2015)
23. Wang, S.I., Manning, C.D.: Fast dropout training. In: *ICML (2)*. pp. 118–126 (2013)
24. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R.S., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. arXiv preprint arXiv:1502.03044 2(3), 5 (2015)
25. Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., Hovy, E.: Hierarchical attention networks for document classification. In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (2016)*
26. Zelikovitz, S., Hirsh, H.: Using lsi for text classification in the presence of background text. In: *Proceedings of the tenth international conference on Information and knowledge management*. pp. 113–118. ACM (2001)
27. Zhang, X., Zhao, J., LeCun, Y.: Character-level convolutional networks for text classification. In: *Advances in Neural Information Processing Systems*. pp. 649–657 (2015)