

Learning text representation using recurrent convolutional neural network with highway layers

Ying Wen^{1,2}, Weinan Zhang¹, Rui Luo¹, and Jun Wang^{1,2}
¹University College London, ²MediaGamma Limited
London, UK
{ying.wen, w.zhang, r.luo, j.wang}@cs.ucl.ac.uk

ABSTRACT

Recently, the rapid development of word embedding and neural networks has brought new inspiration to various NLP and IR tasks. In this paper, we describe a staged hybrid model combining Recurrent Convolutional Neural Networks (RCNN) with highway layers. The highway network module is incorporated in the middle takes the output of the bidirectional Recurrent Neural Network (Bi-RNN) module in the first stage and provides the Convolutional Neural Network (CNN) module in the last stage with the input. The experiment shows that our model outperforms common neural network models (CNN, RNN, Bi-RNN) on a sentiment analysis task. Besides, the analysis of how sequence length influences the RCNN with highway layers shows that our model could learn good representation for the long text.

Keywords

Neural Networks, Highway Networks, Sentiment Analysis

1. INTRODUCTION

Learning good representations for text, such as words, sentences and documents, is essential for information retrieval (IR) and natural language processing (NLP) tasks like text classification and sentiment analysis, which has attracted considerable attention from both academic and industrial communities [2].

It is common to use bag-of-words or bag-of-ngram to represent text [6] and train the models based on such representation. However, each word or n-gram is a unique feature and their interactions and the whole word order in the text are not preserved in the text representation, which limits the functionality of the learning models based on such representation.

In recent years, the word embedding and neural network models have brought new solutions to learn better representations for NLP and IR tasks. The simplest one is Bag-of-Word-Vectors (BOW vector) model. But Landauer et al. [13] estimates that 20% of the meaning of a text comes

from the word order. Therefore, these models are still oversimplified because of the loss of order information. Neural language model [3] was then proposed to leverage word embedding representation to infer the next-word distribution, but it still fails to fully utilize the sequence of the context words.

Recurrent Neural Networks (RNN) can take word order into account, but it suffers from the problem that later words make more influence on the final text representation than former words. However, for sentiment analysis tasks, which is the studied problem of this work, the important words indicating the right sentiment may occur anywhere in the document.

Convolutional Neural Networks (CNN) are naturally capable of solving this problem. CNN treats each words fairly by using the max-pooling layer [9]. Besides, comparing with RNN, which is more natural to process sequence information and get fixed length output, CNN uses sliding windows with different width and filters to perform the feature mapping, then pooling is used to get fixed length output. In addition, CNN is comparatively simple, efficient and has achieved strong empirical performance on the text classification jobs [10]. But CNN also has problems, such as determining the window width and too many filter parameters.

To combine the advantages from Recurrent Neural Network and Convolutional Neural Network, Lai et al. 2015 proposed Recurrent Convolutional Neural Network (RCNN) [12]. This model applies bi-directional recurrent from the beginning to the end of the document to capture the contexts around each word. Then, combine a word and its context to present a word, where concatenate the around contexts and word embedding and filter will be used to calculate the latent semantic vectors. Finally, max-pooling is used to capture the most important factor and make fixed length sentence representation.

In this work, we propose a Recurrent Convolutional Neural Network with Highway Network (RCNN-HW), which incorporates Highway Networks with Recurrent Convolutional Neural Networks. In RCNN-HW model, one highway layer is introduced as intermediate layer between bidirectional RNN and the CNN, which helps to select features individually for each word representation. Therefore, our model further optimises the original RCNN model and achieves better performance.

To summarize, the contributions of our work as follows:

- We first propose the new architecture that incorporates Highway Networks with Recurrent Convolutional Neural Networks.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Neu-IR '16 SIGIR Workshop on Neural Information Retrieval, July 21, 2016, Pisa, Italy

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-2138-9.

DOI: [10.1145/1235](https://doi.org/10.1145/1235)




Figure 1: Architecture of Recurrent Convolutional Neural Network with Highway Layers.

- We apply the proposed model on the sentiment analysis task and achieved superior performance over strong baselines.
- We investigate how sequence length influences the performance of different neural network models, and find RCNN and RCNN-HW models have better performance on longer text.

2. MODEL

In this section, we propose a staged hybrid model combining RCNN with highway layers, which is illustrated in Figure 1. As is shown, the highway network module in the middle takes the output of the bidirectional RNN module in the first stage and provides the CNN module in the last stage with the input.

2.1 Recurrent Neural Networks

It is quite intuitive and straightforward to adopt RNN for learning sequential data due to its nature: RNN is able to process current data instances while taking into account preserved historical information as if it has memory. Among those approaches for enhancement of the long-term memory, Long Short-Term Memory (LSTM) [7] and its variants are

of particular interest due to the great success in practice. Thus we use as the building block of our network a newly proposed variant, Gated Recurrent Unit (GRU) [4], for its simplicity and competitive performance. The expressions of GRU is given as follows:

$$\begin{aligned}
 r_t &= \sigma(\mathbf{W}_r \mathbf{x}_t + \mathbf{U}_r \mathbf{h}_{t-1} + \mathbf{b}_r) \\
 z_t &= \sigma(\mathbf{W}_z \mathbf{x}_t + \mathbf{U}_z \mathbf{h}_{t-1} + \mathbf{b}_z) \\
 \tilde{\mathbf{h}}_t &= \tanh(\mathbf{W}_h \mathbf{x}_t + \mathbf{U}_h (r_t \odot \mathbf{h}_{t-1}) + \mathbf{b}_h) \\
 \mathbf{h}_t &= z_t \odot \mathbf{h}_{t-1} + (1 - z_t) \odot \tilde{\mathbf{h}}_t
 \end{aligned} \tag{1}$$

where \odot represents element-wise multiplication and \mathbf{W} , \mathbf{U} , \mathbf{b} are input weights, recurrent weights and biases, respectively.

We exploit the structure of bidirectional RNN in order to extract both forward-passing and backward-passing information: the context information from the left and right hand side of a particular word. Hence a new context-containing representation $\tilde{\mathbf{x}}_t$ of a word is obtained by simply concatenating the output of the bidirectional RNN ($\mathbf{h}_t^{\leftarrow}, \mathbf{h}_t^{\rightarrow}$) with the word embedding \mathbf{x}_t :

$$\tilde{\mathbf{x}}_t = [\mathbf{h}_t^{\leftarrow} \parallel \mathbf{x}_t \parallel \mathbf{h}_t^{\rightarrow}] \tag{2}$$

2.2 Highway Networks

A one-layer highway network is introduced as the intermediate layer between the bidirectional RNN and the CNN to select features individually for each word representation. The highway layer is expressed as follows:

$$\mathbf{y}_t = \tau \odot g(\mathbf{W}_H \tilde{\mathbf{x}}_t + \mathbf{b}_H) + (1 - \tau) \odot \tilde{\mathbf{x}}_t \tag{3}$$

where g is nonlinear function and $\tau = \sigma(\mathbf{W}_\tau \mathbf{x}_t + \mathbf{b}_\tau)$ represents “transform gate”. Note that the design of highway connection is quite similar with that of GRU’s “update gate” z , which is essentially a variant of “leaky integration” [1]. It allows part of the input information to be carried unchanged to the output while the rest to go through some (nonlinear) transformations. The experiments indicate that this kind of structure will help to extract substantial information.

2.3 Convolutional Neural Networks

In recent years, CNN has achieved great success in CV and has been proved to be effective in various NLP and IR tasks. A CNN architecture is generally composed by a stack of distinct layers mapping input to output via some piece-wise differentiable function. In order to extract new features from the input (phrases/sentences in our case), a convolutional layer essentially applies a set of learnable filters to the input, which have small receptive fields: they are in fact a set of masks with different window sizes h , i.e. they select h consecutive words for the neurons in the layer. For instance, feature \mathbf{c}_i is extracted from a window of word representations with size h : $\mathbf{y}_{i:i+h-1}$ according to the following expression:

$$\mathbf{c}_i = f(\mathbf{w}_{\text{conv}} \cdot \mathbf{y}_{i:i+h-1} + b_{\text{conv}}) \tag{4}$$

where f is a nonlinear function which in our case is the “Rectified Linear Unit” (ReLU), $\mathbf{w}_{\text{conv}}, b_{\text{conv}}$ are the weight and bias.

The filter f will be applied to the word representations of the whole text with length n via a sliding window of size h to establish the feature map:

$$\mathbf{c} = [c_1, c_2, \dots, c_{n-h+1}] \tag{5}$$

Next to the convolution layer, a max-pooling layer is applied to the feature map, which extracts the most significant feature $\hat{c} = \max\{c\}$.

The proposed model utilises multiple filters with window size $h = 1$ to extract features for establishing the representation of the whole text. Note that although it is a common practice to exploit different window sizes, we fix the size to 1 since the bidirectional RNN and highway layer in the earlier layers have already obtained the context information around each word and thus we are free from applying filters in different window sizes.

3. EXPERIMENT

To demonstrate the effectiveness of the proposed model, we perform the experiment on IMDB dataset [14] which contains 25,000 movie reviews for training. Each review is labelled by human with 1 represents positive sentiment and 0 represents negative sentiment. The review in the dataset has 267.9 words in average and standard deviation of review length is 198.8. To test the performance of learned text presentation of models, we run sentiment prediction on this dataset and report the accuracy evaluation metric. Besides, we are particularly interested in the relationship between input sequence length and performance of the model, and set up the experiment to test the performance of the model over different input sequence lengths.

We compare our model RCNN with highway layers (RCNN-HW) with the following baseline neural network models for document level sentiment prediction:

- **Sum-of-Word-Vector (COW)**: simply sums up all the word embeddings as text representation.
- **LSTM**: takes the average of the LSTM’s hidden states of all words is used as text representation [7].
- **Bi-LSTM**: similar to LSTM but exploits bidirectional LSTM [16] on the sequence to get text representation.
- **CNN**: performs 1-dimensional convolution followed by 1-dimensional max-pooling with multiple filters on the sequence [9, 10].
- **CNN+LSTM**: combines CNN and RNN (LSTM) by using CNN’s output as LSTM’s input.
- **RCNN**: uses bidirectional RNN’s output as CNN’s input [12].

We use neural network models above to learn the text representation and adopt softmax output to make prediction. For the neural network models, hyperparameter tuning has great influence on its performance. For all the models we have done grid search with reasonable hyperparameters and tried different stochastic gradient descent training methods (e.g. adam, RMSprop, adadelta) [11]. For our model, RCNN with highway layers¹, we choose ‘RMSprop’ as training method with batch size of 32, the hidden layer dimension of 32, and the filter number of 256. Besides, we also tested the performance of different neural network models with multiple input sequence lengths, the RCNN based models get better performance with longer sequence length, in our

¹The code is available in https://github.com/wenyng45/deep_learning_tutorial/tree/master/rcnn-hw

Table 1: Accuracy for different neural models on IMDB dataset

Model	Accuracy
COW	0.890
LSTM	0.885
Bi-LSTM	0.881
CNN+LSTM	0.890
CNN	0.895
RCNN	0.900
RCNN-HW	0.903

Table 2: Accuracy for RCNN models with/without highway layers

RCNN based model	Accuracy
Without Highway Layers	0.900
One Highway Layers	0.903
Two Highway Layers	0.903
One MLP Layer	0.899

experiment, we choose the best performance of the model over different sequence lengths. Note, we did not apply pre-training methods like using pre-trained word2vec [15], and regularizer such as Dropout [17] which may bring further improvement [10].

4. RESULTS AND DISCUSSION

The experiment results can be found in Table 1. We compare our RCNN with highway layers with original RCNN model and find that the performance with highway layers are always better than those without. To quantitatively investigate the effect of highway network, we set up an experiment with RCNN based models. We train the RCNN model without highway layers, with one layer highway, two layers highway and one layer MLP respectively with the same parameters, then compare their accuracy. The results are shown in Table 2. We can find that having one to two highway layers is important, but more highway layers do not improve the performance. Besides, one MLP layer does not gain more improvement than highway layers. We think the reason why the highway layer works is because the highway layer helps select the features of word representation.

Besides, we also compare with CNN, RNN and COW models. Empirical experiments show that CNN models usually yield better performance than RNNs which include LSTM and Bi-LSTM. It may be caused by RNN based models, especially the LSTMs, are hard to be trained, these models are really sensitive to the hyperparameters and latter words make more influence on the final text representation than former words in RNN models. CNN models usually perform remarkably well on many NLP and IR tasks [8, 9, 10, 18]. CNN based models use a fixed window of words as contextual information and the performance of a CNN is influenced by the window size. A small window may result in a loss of some long-distance patterns, whereas large windows will lead to data sparsity [5, 12]. However, RCNNs outperform the CNN and the CNN-LSTM model, we know that convolution structure can capture local context information, and recurrent can capture global information. We consider that in CNN-LSTM model, it does not make sense to use convolution layer and max-pooling layer before recur-

rent layer. Such architecture means using local information features extracted from CNN as RNN's input, but local information features do not have sequence relationship. As for why RCNN based models outperform CNN, we think our context-containing representation of a word obtained by simply concatenating the output of the bidirectional RNN with the word embedding which has encoded context information would be a better input to the CNN comparing with original word embedding.




Figure 2: Accuracy curve for how sequence length influences the performance of different neural network models

We further investigate how the input sequence length influences the performance of different neural network models. The average text length in our experiment distaste is 267 and the standard deviation of the text length is around 200. Therefore, we consider the input sequence length from 100 to 500 with step 100 to train the neural network models. The prediction accuracy for different neural network models with various input sequence length are shown in Figure 2. In the figure, we can observe that when the sequence length is less than 200, all models get bad performance since the short input length results in information loss. As the input sequence length increases, models receive longer input sequence which may bring more useful information as well as more noise information. From the figure, we can find that RCNN based models, especially our RCNN-HW model, obtain much higher improvement from longer input sequences. Meanwhile, the other models get lower improvement and even get worse performance than shorter input sequence length. We think this is because RCNN based models can preserve longer contextual information thanks to their recurrent structure and introduce less noise for their convolution structure with max-pooling. And our RCNN-HW can further reduce the noise with a highway layer to perform feature selection, which brings the best performance in longer input sequence lengths.

We can take an example to illustrate why our RCNN-HW model outperforms than others. There is a positive review “I borrowed this movie despite its extremely low rating, because I wanted to see how the crew manages to animate the presence of multiple worlds. As a matter of fact, they didn’t

- at least, so it seems ... But I closed my eyes. When I opened them again - he was gone.” in the dataset, and its length is 498. This review can be easily classified as negative thus the double negative and the long length. Simple model can hard to capture this type of structure and handle the long-term dependency, so as COW, CNN, RNN models, they predict this sample as a negative review. But our RCNN-HW model can deal with this situation and makes right predictbbbbion.

5. CONCLUSION

We have introduced a recurrent convolutional neural network with highway layers to learn text representation for sentiment analysis. The experiment demonstrates that our model RCNN-HW not only outperforms CNN and RNN models, but also works better than the RCNN without highway layer. Besides, since our model yield better performance on longer text, it would be interesting to see if the model introduced in this paper works well in learning the text representation for the other NLP and IR tasks with relatively long documents.

6. ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their helpful feedback and suggestions.

7. REFERENCES

- [1] Y. Bengio, N. Boulanger-Lewandowski, and R. Pascanu. Advances in optimizing recurrent networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 8624–8628. IEEE, 2013.
- [2] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):1798–1828, 2013.
- [3] Y. Bengio, H. Schwenk, J.-S. Senécal, F. Morin, and J.-L. Gauvain. Neural probabilistic language models. In *Innovations in Machine Learning*, pages 137–186. Springer, 2006.
- [4] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [5] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [6] W. B. Croft, D. Metzler, and T. Strohman. *Search engines: Information retrieval in practice*, volume 283. Addison-Wesley Reading, 2010.
- [7] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [8] B. Hu, Z. Lu, H. Li, and Q. Chen. Convolutional neural network architectures for matching natural language sentences. In *Advances in Neural Information Processing Systems*, pages 2042–2050, 2014.
- [9] N. Kalchbrenner, E. Grefenstette, and P. Blunsom. A convolutional neural network for modelling sentences. In *ACL*, 2014.

- [10] Y. Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.
- [11] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [12] S. Lai, L. Xu, K. Liu, and J. Zhao. Recurrent convolutional neural networks for text classification. In *AAAI*, pages 2267–2273, 2015.
- [13] T. K. Landauer. On the computational basis of learning and cognition: Arguments from lsa. *Psychology of learning and motivation*, 41:43–84, 2002.
- [14] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 142–150. Association for Computational Linguistics, 2011.
- [15] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc., 2013.
- [16] M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *Signal Processing, IEEE Transactions on*, 45(11):2673–2681, 1997.
- [17] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [18] Y. Zhang and B. Wallace. A sensitivity analysis of (and practitioners’ guide to) convolutional neural networks for sentence classification. *arXiv preprint arXiv:1510.03820*, 2015.