

# Convolutional Neural Networks with Recurrent Neural Filters

Yi Yang

Bloomberg

New York, NY 10022

yyang464@bloomberg.net

## Abstract

We introduce a class of convolutional neural networks (CNNs) that utilize recurrent neural networks (RNNs) as convolution filters. A convolution filter is typically implemented as a linear affine transformation followed by a non-linear function, which fails to account for language compositionality. As a result, it limits the use of high-order filters that are often warranted for natural language processing tasks. In this work, we model convolution filters with RNNs that naturally capture compositionality and long-term dependencies in language. We show that simple CNN architectures equipped with recurrent neural filters (RNFs) achieve results that are on par with the best published ones on the Stanford Sentiment Treebank and two answer sentence selection datasets.<sup>1</sup>

## 1 Introduction

Convolutional neural networks (CNNs) have been shown to achieve state-of-the-art results on various natural language processing (NLP) tasks, such as sentence classification (Kim, 2014), question answering (Dong et al., 2015), and machine translation (Gehring et al., 2017). In an NLP system, a convolution operation is typically a sliding window function that applies a convolution filter to every possible window of words in a sentence. Hence, the key components of CNNs are a set of convolution filters that compose low-level word features into higher-level representations.

Convolution filters are usually realized as linear systems, as their outputs are affine transformations of the inputs followed by some non-linear activation functions. Prior work directly adopts a linear convolution filter to NLP problems by utilizing a concatenation of embeddings of a window of words as the input, which leverages word order information in a shallow additive way. As early

<sup>1</sup>The code is available at <https://github.com/bloomberg/cnn-rnf>.

CNN architectures have mainly drawn inspiration from models of the primate visual system, the necessity of coping with language compositionality is largely overlooked in these systems. Due to the linear nature of the convolution filters, they lack the ability to capture complex language phenomena, such as compositionality and long-term dependencies.

To overcome this, we propose to employ recurrent neural networks (RNNs) as convolution filters of CNN systems for various NLP tasks. Our recurrent neural filters (RNFs) can naturally deal with language compositionality with a recurrent function that models word relations, and they are also able to implicitly model long-term dependencies. RNFs are typically applied to word sequences of moderate lengths, which alleviates some well-known drawbacks of RNNs, including their vulnerability to the gradient vanishing and exploding problems (Pascanu et al., 2013). As in conventional CNNs, the computation of the convolution operation with RNFs can be easily parallelized. As a result, RNF-based CNN models can be 3-8x faster than their RNN counterparts.

We present two RNF-based CNN architectures for sentence classification and answer sentence selection problems. Experimental results on the Stanford Sentiment Treebank and the QASent and WikiQA datasets demonstrate that RNFs significantly improve CNN performance over linear filters by 4-5% accuracies and 3-6% MAP scores respectively. Analysis results suggest that RNFs perform much better than linear filters in detecting longer key phrases, which provide stronger cues for categorizing the sentences.

## 2 Approach

The aim of a convolution filter is to produce a local feature for a window of words. We describe a novel approach to learning filters using RNNs, which is especially suitable for NLP problems.

We then present two CNN architectures equipped with RNFs for sentence classification and sentence matching tasks respectively.

## 2.1 Recurrent neural filters

We briefly review the linear convolution filter implementation by Kim (2014), which has been widely adopted in later works. Consider an  $m$ -gram word sequence  $[\mathbf{x}_i, \dots, \mathbf{x}_{i+m-1}]$ , where  $\mathbf{x}_i \in \mathbb{R}^k$  is a  $k$ -dimensional word vector. A linear convolution filter is a function applied to the  $m$ -gram to produce a feature  $c_{i,j}$ :

$$c_{i,j} = f(\mathbf{w}_j^\top \mathbf{x}_{i:i+m-1} + b_j), \quad (1)$$

$$\mathbf{x}_{i:i+m-1} = \mathbf{x}_i \oplus \mathbf{x}_{i+1} \oplus \dots \oplus \mathbf{x}_{i+m-1},$$

where  $\oplus$  is the concatenation operator,  $b_j$  is a bias term, and  $f$  is a non-linear activation function. We typically use multiple independent linear filters to yield a feature vector  $\mathbf{c}_i$  for the word sequence. Linear convolution filters make strong assumptions about language that could harm the performance of NLP systems. First, linear filters assume local compositionality and ignore long-term dependencies in language. Second, they use separate parameters for each value of the time index, which hinders parameter sharing for the same word type (Goodfellow et al., 2016). The assumptions become more problematic if we increase the window size  $m$ .

We propose to address the limitations by employing RNNs to realize convolution filters, which we term *recurrent neural filters* (RNFs). RNFs compose the words of the  $m$ -gram from left to right using the same recurrent unit:

$$\mathbf{h}_t = \text{RNN}(\mathbf{h}_{t-1}, \mathbf{x}_t), \quad (2)$$

where  $\mathbf{h}_t$  is a hidden state vector that encoded information about previously processed words, and the function RNN is a recurrent unit such as Long Short-Term Memory (LSTM) unit (Hochreiter and Schmidhuber, 1997) or Gated Recurrent Unit (GRU) (Cho et al., 2014). We use the last hidden state  $\mathbf{h}_{i+m-1}$  as the RNF output feature vector  $\mathbf{c}_i$ . Features learned by RNFs are interdependent of each other, which permits the learning of complementary information about the word sequence. The left-to-right word composing procedure in RNFs preserves word order information and implicitly models long-term dependencies in language. RNFs can be treated as simple drop-in replacements of linear filters and potentially adopted in numerous CNN architectures.

## 2.2 CNN architectures

**Sentence encoder** We use a CNN architecture with one convolution layer followed by one max pooling layer to encode a sentence. In particular, the RNFs are applied to every possible window of  $m$  words in the sentence  $\{\mathbf{x}_{1:m}, \mathbf{x}_{2:m+1}, \dots, \mathbf{x}_{n-m+1:n}\}$  to generate feature maps  $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_{n-m+1}]$ . Then a max-over-time pooling operation (Collobert et al., 2011) is used to produce a fixed size sentence representation:  $\mathbf{v} = \max\{\mathbf{C}\}$ .

**Sentence classification** We use a CNN architecture that is similar to the CNN-static model described in Kim (2014) for sentence classification. After obtaining the sentence representation  $\mathbf{v}$ , a fully connected softmax layer is used to map  $\mathbf{v}$  to an output probability distribution. The network is optimized against a cross-entropy loss function.

**Sentence matching** We exploit a CNN architecture that is nearly identical to the CNN-Cnt model introduced by Yang et al. (2015). Let  $\mathbf{v}_1$  and  $\mathbf{v}_2$  be the vector representations of the two sentences. A bilinear function is applied to  $\mathbf{v}_1$  and  $\mathbf{v}_2$  to produce a sentence matching score. The score is combined with two word matching count features and fed into a sigmoid layer. The output of the sigmoid layer is used by binary cross-entropy loss to optimize the model.

## 3 Experiments

We evaluate RNFs on some of the most popular datasets for the sentence classification and sentence matching tasks. After describing the experimental setup, we compare RNFs against both linear filters and conventional RNN models, and report our findings in § 3.2.

### 3.1 Experimental settings

**Data** We use the Stanford Sentiment Treebank (Socher et al., 2013) in our sentence classification experiments. We report accuracy results for both binary classification and fine-grained classification settings. Two answer sentence selection datasets, QASent (Wang et al., 2007) and WikiQA (Yang et al., 2015), are adopted in our sentence matching experiments. We use MAP and MRR to evaluate the performance of answer sentence selection models.

**Competitive systems** We consider CNN variants with linear filters and RNFs. For RNFs, we

adopt two implementations based on GRUs and LSTMs respectively. We also compare against the following RNN variants: GRU, LSTM, GRU with max pooling, and LSTM with max pooling.<sup>2</sup> We use the publicly available 300-dimensional GloVe vectors (Pennington et al., 2014) pre-trained with 840 billion tokens to initialize the word embeddings for all the models. The word vectors are fixed during downstream training. Finally, we report best published results for each dataset.<sup>3</sup>

**Parameter tuning** For all the experiments, we tune hyperparameters on the development sets and report results obtained with the selected hyperparameters on the test sets. After the preliminary search, we choose the number of hidden units (feature maps) from {200, 300, 400}, and the filter window width from {2, 3, 4, 5} for linear filters and from {5, 6, 7, 8} for RNFs. Linear filters tend to perform well with small window widths, while RNFs work better with larger window widths. We apply dropout to embedding layers, pooling layers, RNN input layers, and RNN recurrent layers. The dropout rates are selected from {0, 0.2, 0.4}, where 0 indicates that dropout is not used for the specific layer. Optimization is performed by Adam (Kingma and Ba, 2015) with early stopping. We conduct random search with a budget of 100 runs to seek the best hyperparameter configuration for each system.

### 3.2 Results

The evaluation results on sentiment classification and answer sentence selection are shown in Table 1 and Table 2 respectively. RNFs significantly outperform linear filters on both tasks. In fact, we find that CNN-RNF variants significantly outperform CNN-linear-filter on nearly every hyperparameter configuration in our experiments. On the Stanford Sentiment Treebank, CNN-RNF-LSTM outperforms CNN-linear-filter by 5.4% and 3.9% accuracies on the fine-grained and binary classification settings respectively. We observe similar performance boosts by adopting RNFs for CNNs on the QASent and WikiQA test sets. CNN-RNF-GRU improves upon CNN-linear-filter by 3.7% MRR score on QASent and CNN-RNF-LSTM performs better than CNN-linear-filter by 6.1%

<sup>2</sup>Max pooling performed better than mean pooling in our preliminary experiments.

<sup>3</sup>We exclude results obtained from systems using external resources beyond word embeddings.

System	Fine-grained	Binary
<i>CNN variants</i>		
CNN-linear-filter	48.0	86.1
CNN-RNF-GRU	53.0	<b>90.0</b>
CNN-RNF-LSTM	<b>53.4</b>	<b>90.0</b>
<i>RNN variants</i>		
GRU	50.5	88.7
LSTM	50.3	89.3
GRU-maxpool	51.7	89.7
LSTM-maxpool	51.6	89.8
<i>Best published results</i>		
(Lei et al., 2017)	<u>53.2</u>	<u>89.9</u>

Table 1: Accuracy results on the Stanford Sentiment Treebank. The best results obtained from our implementations are in **bold**. The best published results are underlined.

MAP score on WikiQA. In particular, CNN-RNF-LSTM achieves 53.4% and 90.0% accuracies on the fine-grained and binary sentiment classification tasks respectively, which match the state-of-the-art results on the Stanford Sentiment Treebank. CNN-RNF-LSTM also obtains competitive results on answer sentence selection datasets, despite the simple model architecture compared to state-of-the-art systems.

Conventional RNN models clearly benefit from max pooling, especially on the task of answer sentence selection. Like RNF-based CNN models, max-pooled RNNs also consist of two essential layers. The recurrent layer learns a set of hidden states corresponding to different time steps, and the max pooling layer extracts the most salient information from the hidden states. However, a hidden state in RNNs encodes information about all the previous time steps, while RNFs focus on detecting local features from a window of words that can be more relevant to specific tasks. As a result, RNF-based CNN models perform consistently better than max-pooled RNN models.

CNN-RNF models are much faster to train than their corresponding RNN counterparts, despite they have the same numbers of parameters, as RNFs are applied to word sequences of shorter lengths and the computation is parallelizable. The training of CNN-RNF-LSTM models takes 10-20 mins on the Stanford Sentiment Treebank, which is 3-8x faster than the training of LSTM models, on an NVIDIA Tesla P100 GPU accelerator.

System	QASent		WikiQA	
	MAP	MRR	MAP	MRR
<i>CNN variants</i>				
CNN-linear-filter	0.750	0.813	0.668	0.687
CNN-RNF-GRU	0.773	<b>0.850</b>	0.726	0.738
CNN-RNF-LSTM	<b>0.780</b>	0.841	<b>0.729</b>	<b>0.747</b>
<i>RNN variants</i>				
GRU	0.739	0.815	0.658	0.678
LSTM	0.745	0.806	0.651	0.671
GRU-maxpool	0.744	0.814	0.712	0.724
LSTM-maxpool	0.762	0.825	0.701	0.711
<i>Best published results</i>				
(Rao et al., 2016)	<u>0.801</u>	<u>0.877</u>	0.701	0.718
(Wang and Jiang, 2017)	-	-	<u>0.743</u>	<u>0.755</u>

Table 2: Evaluation results on answer sentence selection datasets. The best results obtained from our implementations are in **bold**. The best published results are underlined.




Figure 1: Local label consistency ratio results computed on SST dev set.

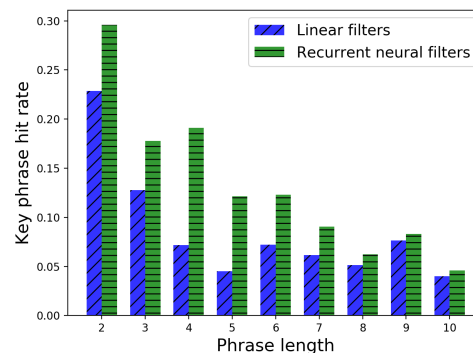


Figure 2: Key phrase hit rate results computed on SST dev set. The RNFs are implemented with LSTMs.

## 4 Analysis

We now investigate why RNFs are more effective than linear convolution filters on the binary sentiment classification task. We perform quantitative analysis on the development set of the Stanford Sentiment Treebank (SST), in which sentiment labels for some phrases are also available.

**Local label consistency (LLC) ratio** We first inspect whether longer phrases have a higher tendency to express the same sentiment as the entire sentence. We define the *local label consistency (LLC) ratio* as the ratio of  $m$ -grams that share the same sentiment labels as the original sentences. The LLC ratios with respect to different phrase lengths are shown in Figure 1. Longer phrases are

more likely to convey the same sentiments as the original sentences. Therefore, the ability to model long phrases is crucial to convolution filters.

**Key phrase hit rate** We examine linear filters and RNFs on the ability to detect a key phrase in a sentence. Specifically, we define the key phrases for a sentence to be the set of phrases that are labeled with the same sentiment as the original sentence. The *key phrase hit rate* is then defined as the ratio of filter-detected  $m$ -grams that fall into the corresponding key phrase sets. The filter-detected  $m$ -gram of a sentence is the one whose convolution feature vector has the shortest Euclidean distance to the max-pooled vector. The hit rate results computed on SST dev set are presented in Figure 2. As shown, RNFs consistently per-

form better than linear filters on identifying key phrases across different phrase lengths, especially on phrases of moderate lengths (4-7). The results suggest that RNFs are superior to linear filters, as they can better model longer phrases whose labels are more consistent with the sentences.<sup>4</sup>

## 5 Related work

Linear convolution filters are dominated in CNN-based systems for both computer vision and natural language processing tasks. One exception is the work of [Zoumpourlis et al. \(2017\)](#), which proposes a convolution filter that is a function with quadratic forms through the Volterra kernels. However, this non-linear convolution filter is developed in the context of a computational model of the visual cortex, which is not suitable for NLP problems. In contrast, RNFs are specifically derived for NLP tasks, in which a different form of non-linearity, language compositionality, often plays a critical role.

Several works have employed neural network architectures that contain both CNN and RNN components to tackle NLP problems. [Tan et al. \(2016\)](#) present a deep neural network for answer sentence selection, in which a convolution layer is applied to the output of a BiLSTM layer for extracting sentence representations. [Ma and Hovy \(2016\)](#) propose to compose character representations of a word using a CNN, whose output is then fed into a BiLSTM for sequence tagging. [Kalchbrenner and Blunsom \(2013\)](#) introduce a neural architecture that uses a sentence model based on CNNs and a discourse model based on RNNs. Their system achieves state-of-the-art results on the task of dialogue act classification. Instead of treating an RNN and a CNN as isolated components, our work directly marries RNNs with the convolution operation, which illustrates a new direction in mixing RNNs with CNNs.

## 6 Conclusion and future work

We present RNFs, a new class of convolution filters based on recurrent neural networks. RNFs sequentially apply the same recurrent unit to words of a phrase, which naturally capture language compositionality and implicitly model long-term

<sup>4</sup>Phrases of very long lengths (8-10) are rarely annotated in SST dev data, which could explain why linear filters and RNFs achieve similar hit rates, as small data sample often leads to high variance.

dependencies. Experiments on sentiment classification and answer sentence selection tasks demonstrate that RNFs give a significant boost in performance compared to linear convolution filters. RNF-based CNNs also outperform a variety of RNN-based models, as they focus on extracting local information that could be more relevant to particular problems. The quantitative analysis reveals that RNFs can handle long phrases much better than linear filters, which explains their superiority over the linear counterparts. In the future, we would like to investigate the effectiveness of RNFs on a wider range of NLP tasks, such as natural language inference and machine translation.

## 7 Acknowledgments

We thank Chunyang Xiao for helping us to run the analysis experiments. We thank Kazi Shefaet Rahman, Ozan Irsoy, Chen-Tse Tsai, and Lingjia Deng for their valuable comments on earlier versions of this paper. We also thank the EMNLP reviewers for their helpful feedback.

## References

- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Li Dong, Furu Wei, Ming Zhou, and Ke Xu. 2015. Question answering over freebase with multi-column convolutional neural networks. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. 2016. *Deep learning*, volume 1. MIT press Cambridge.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8).



- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent convolutional neural networks for discourse compositionality. In *Proceedings of the Workshop on Continuous Vector Space Models and their Compositionality*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Tao Lei, Wengong Jin, Regina Barzilay, and Tommi Jaakkola. 2017. Deriving neural architectures from sequence and graph kernels. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Jinfeng Rao, Hua He, and Jimmy Lin. 2016. Noise-contrastive estimation for answer selection with deep neural networks. In *Proceedings of the ACM International on Conference on Information and Knowledge Management (CIKM)*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Ming Tan, Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2016. Lstm-based deep learning models for non-factoid answer selection. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Mengqiu Wang, Noah A Smith, and Teruko Mitamura. 2007. What is the jeopardy model? a quasi-synchronous grammar for qa. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Shuohang Wang and Jing Jiang. 2017. A compare-aggregate model for matching text sequences. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *Proceedings of Empirical Methods for Natural Language Processing (EMNLP)*.
- Georgios Zoumpourlis, Alexandros Doumanoglou, Nicholas Vretos, and Petros Daras. 2017. Non-linear convolution filters for cnn-based learning. In *Proceedings of the International Conference on Computer Vision (ICCV)*.