
Supervised and Semi-Supervised Text Categorization using LSTM for Region Embeddings

Rie Johnson

RJ Research Consulting, Tarrytown NY, USA

RIEJOHNSON@GMAIL.COM

Tong Zhang

Big Data Lab, Baidu Inc, Beijing, China

TONGZHANG@BAIDU.COM

Abstract

One-hot CNN (convolutional neural network) has been shown to be effective for text categorization (Johnson & Zhang, 2015a;b). We view it as a special case of a general framework which jointly trains a linear model with a non-linear feature generator consisting of ‘*text region embedding + pooling*’. Under this framework, we explore a more sophisticated region embedding method using *Long Short-Term Memory (LSTM)*. LSTM can embed text regions of variable (and possibly large) sizes, whereas the region size needs to be fixed in a CNN. We seek effective and efficient use of LSTM for this purpose in the supervised and semi-supervised settings. The best results were obtained by combining region embeddings in the form of LSTM and convolution layers trained on unlabeled data. The results indicate that on this task, embeddings of text regions, which can convey complex concepts, are more useful than embeddings of single words in isolation. We report performances exceeding the previous best results on four benchmark datasets.

1. Introduction

Text categorization is the task of assigning labels to documents written in a natural language, and it has numerous real-world applications including sentiment analysis as well as traditional topic assignment tasks. The state-of-the-art methods for text categorization had long been linear predictors (e.g., SVM with a linear kernel) with either bag-of-word or bag-of- n -gram vectors (hereafter *bow*) as input, e.g., (Joachims, 1998; Lewis et al., 2004). This, however,

has changed recently. Non-linear methods that can make effective use of word order have been shown to produce more accurate predictors than the traditional bow-based linear models, e.g., (Dai & Le, 2015; Zhang et al., 2015). In particular, let us first focus on *one-hot CNN* which we proposed in JZ15 (Johnson & Zhang, 2015a;b).

A *convolutional neural network (CNN)* (LeCun et al., 1986) is a feedforward neural network with convolution layers interleaved with pooling layers, originally developed for image processing. In its convolution layer, a small region of data (e.g., a small square of image) at every location is converted to a low-dimensional vector with information relevant to the task being preserved, which we loosely term ‘*embedding*’. The *embedding function* is shared among all the locations, so that useful features can be detected irrespective of their locations. In its simplest form, one-hot CNN works as follows. A document is represented as a sequence of *one-hot vectors* (each of which indicates a word by the position of a 1); a convolution layer converts small regions of the document (e.g., “I love it”) to low-dimensional vectors at every location (*embedding of text regions*); a pooling layer aggregates the region embedding results to a document vector by taking component-wise maximum or average; and the top layer classifies a document vector with a linear model (Figure 1). The one-hot CNN and its semi-supervised extension were shown to be superior to a number of previous methods.

In this work, we consider a more general framework (subsuming one-hot CNN) which *jointly trains a feature generator and a linear model*, where the feature generator consists of ‘*region embedding + pooling*’. The specific region embedding function of one-hot CNN takes the simple form

$$\mathbf{v}(\mathbf{x}_\ell) = \max(0, \mathbf{W}\mathbf{x}_\ell + \mathbf{b}), \quad (1)$$

where \mathbf{x}_ℓ is a concatenation of one-hot vectors (therefore, ‘one-hot’ in the name) of the words in the ℓ -th region (of a fixed size), and the weight matrix \mathbf{W} and the bias vector \mathbf{b} need to be trained. It is simple and fast to compute, and

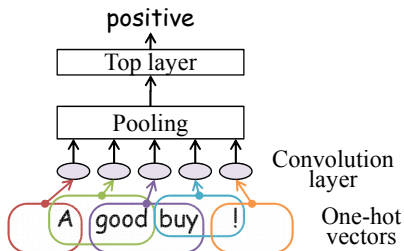


Figure 1. One-hot CNN (*oh-CNN*) [JZ15a]




Figure 2. Word vector LSTM (*wv-LSTM*) as in [DL15].




Figure 3. One-hot LSTM with pooling (*oh-LSTMp*).

considering its simplicity, the method works surprisingly well if the region size is appropriately set. However, there are also potential shortcomings. The region size must be fixed, which may not be optimal as the size of relevant regions may vary. Practically, the region size cannot be very large as the number of parameters to be learned (components of \mathbf{W}) depends on it. JZ15 proposed variations to alleviate these issues. For example, a bow-input variation allows \mathbf{x}_ℓ above to be a bow vector of the region. This enables a larger region, but at the expense of losing word order in the region and so its use may be limited.

In this work, we build on the general framework of ‘region embedding + pooling’ and explore a more sophisticated region embedding via *Long Short-Term Memory (LSTM)*, seeking to overcome the shortcomings above, in the supervised and semi-supervised settings. LSTM (Hochreiter & Schmidhuder, 1997) is a recurrent neural network. In its typical applications to text, an LSTM takes words in a sequence one by one; i.e., at time t , it takes as input the t -th word and the output from time $t - 1$. Therefore, the output from each time step can be regarded as the embedding of the sequence of words that have been seen so far (or a relevant part of it). It is designed to enable learning of dependencies over larger time lags than feasible with traditional recurrent networks. That is, an LSTM can be used to embed text regions of variable (and possibly large) sizes.

We pursue the best use of LSTM for our purpose, and then compare the resulting model with the previous best methods including one-hot CNN and previous LSTM. Our strategy is to *simplify the model* as much as possible, including elimination of a word embedding layer routinely used to produce input to LSTM. Our findings are three-fold. First, in the supervised setting, our simplification strategy leads to higher accuracy and faster training than previous LSTM. Second, accuracy can be further improved by training LSTMs on unlabeled data for learning useful *region embeddings* and using them to produce *additional input*. Third, both our LSTM models and one-hot CNN strongly outperform other methods including previous LSTM. The best results are obtained by combining the two types of region embeddings (*LSTM embed-*

dings and *CNN embeddings*) trained on unlabeled data, indicating that their strengths are complementary. Overall, our results show that for text categorization, embeddings of text regions, which can convey higher-level concepts than single words in isolation, are useful, and that useful region embeddings can be learned without going through word embedding learning. We report performances exceeding the previous best results on four benchmark datasets. Our code and experimental details are available at http://riejohnson.com/cnn_download.html.

1.1. Preliminary

On text, LSTM has been used for labeling or generating words. It has been also used for representing short sentences mostly for sentiment analysis, and some of them rely on syntactic parse trees; see e.g., (Zhu et al., 2015; Tang et al., 2015; Tai et al., 2015; Le & Zuidema, 2015). Unlike these studies, this work as well as JZ15 focuses on classifying general full-length documents without any special linguistic knowledge. Similarly, DL15 (Dai & Le, 2015) applied LSTM to categorizing general full-length documents. Therefore, our empirical comparisons will focus on DL15 and JZ15, both of which reported new state of the art results. Let us first introduce the general LSTM formulation, and then briefly describe DL15’s model as it illustrates the challenges in using LSTMs for this task.

LSTM While several variations exist, we base our work on the following LSTM formulation, which was used in, e.g., (Zaremba & Sutskever, 2014)

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{W}^{(i)}\mathbf{x}_t + \mathbf{U}^{(i)}\mathbf{h}_{t-1} + \mathbf{b}^{(i)}), \\ \mathbf{o}_t &= \sigma(\mathbf{W}^{(o)}\mathbf{x}_t + \mathbf{U}^{(o)}\mathbf{h}_{t-1} + \mathbf{b}^{(o)}), \\ \mathbf{f}_t &= \sigma(\mathbf{W}^{(f)}\mathbf{x}_t + \mathbf{U}^{(f)}\mathbf{h}_{t-1} + \mathbf{b}^{(f)}), \\ \mathbf{u}_t &= \tanh(\mathbf{W}^{(u)}\mathbf{x}_t + \mathbf{U}^{(u)}\mathbf{h}_{t-1} + \mathbf{b}^{(u)}), \\ \mathbf{c}_t &= \mathbf{i}_t \odot \mathbf{u}_t + \mathbf{f}_t \odot \mathbf{c}_{t-1}, \\ \mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t), \end{aligned}$$

where \odot denotes element-wise multiplication and σ is an element-wise squash function to make the gating values in $[0, 1]$. We fix σ to sigmoid. $\mathbf{x}_t \in \mathbb{R}^d$ is the input from

the lower layer at time step t , where d would be, for example, size of vocabulary if the input was a one-hot vector representing a word, or the dimensionality of word vector if the lower layer was a word embedding layer. With q LSTM units, the dimensionality of the weight matrices and bias vectors, which need to be trained, are $\mathbf{W}^{(\cdot)} \in \mathbb{R}^{q \times d}$, $\mathbf{U}^{(\cdot)} \in \mathbb{R}^{q \times q}$, and $\mathbf{b}^{(\cdot)} \in \mathbb{R}^q$ for all types (i, o, f, u) . The centerpiece of LSTM is the *memory cells* \mathbf{c}_t , designed to counteract the risk of vanishing/exploding gradients, thus enabling learning of dependencies over larger time lags than feasible with traditional recurrent networks. The *forget gate* \mathbf{f}_t (Gers et al., 2000) is for resetting the memory cells. The *input gate* \mathbf{i}_t and *output gate* \mathbf{o}_t control the input and output of the memory cells.

Word-vector LSTM (wv-LSTM) [DL15] DL15’s application of LSTM to text categorization is straightforward. As illustrated in Figure 2, for each document, the output of the LSTM layer is the output of the last time step (corresponding to the last word of the document), which represents the whole document (document embedding). Like many other studies of LSTM on text, words are first converted to low-dimensional dense word vectors via a word embedding layer; therefore, we call it *word-vector LSTM* or *wv-LSTM*. DL15 observed that wv-LSTM underperformed linear predictors and its training was unstable. This was attributed to the fact that documents are long.

In addition, we found that training and testing of wv-LSTM is time/resource consuming. To put it into perspective, using a GPU, one epoch of wv-LSTM training takes nearly 20 times longer than that of one-hot CNN training even though it achieves poorer accuracy (the first two rows of Table 1). This is due to the sequential nature of LSTM, i.e., computation at time t requires the output of time $t - 1$, whereas modern computation depends on parallelization for speed-up. Documents in a mini-batch can be processed in parallel, but the variability of document lengths reduces the degree of parallelization¹.

It was shown in DL15 that training becomes stable and accuracy improves drastically when LSTM and the word embedding layer are jointly *pre-trained* with either the language model learning objective (predicting the next word) or autoencoder objective (memorizing the document).

2. Supervised LSTM for text categorization

Within the framework of ‘region embedding + pooling’ for text categorization, we seek effective and efficient use of LSTM as an alternative region embedding method. This

¹ (Sutskever et al., 2014) suggested making each mini-batch consist of sequences of similar lengths, but we found that on our tasks this strategy slows down convergence presumably by hampering the stochastic nature of SGD.

section focuses on an end-to-end supervised setting so that there is no additional data (e.g., unlabeled data) or additional algorithm (e.g., for learning a word embedding). Our general strategy is to simplify the model as much as possible. We start with elimination of the word embedding layer so that one-hot vectors are directly fed to LSTM, which we call *one-hot LSTM* in short.

2.1. Elimination of the word embedding layer

Facts: A word embedding is a linear operation that can be written as $\mathbf{V}\mathbf{x}_t$ with \mathbf{x}_t being a one-hot vector and columns of \mathbf{V} being word vectors. Therefore, by replacing the LSTM weights $\mathbf{W}^{(\cdot)}$ with $\mathbf{W}^{(\cdot)}\mathbf{V}$ and removing the word embedding layer, a word-vector LSTM can be turned into a one-hot LSTM without changing the model behavior. Thus, word-vector LSTM is not more expressive than one-hot LSTM; rather, a merit, if any, of training with a word embedding layer would be through imposing restrictions (e.g., a low-rank \mathbf{V} makes a less expressive model) to achieve good prior/regularization effects.

In the end-to-end supervised setting, a word embedding matrix \mathbf{V} would need to be initialized randomly and trained as part of the model. In the preliminary experiments under our framework, we were unable to improve accuracy over one-hot LSTM by inclusion of such a randomly initialized word embedding layer; i.e., random vectors failed to provide good prior effects. Instead, demerits were evident – more meta-parameters to tune, poor accuracy with low-dimensional word vectors, and slow training/testing with high-dimensional word vectors as they are dense.

If a word embedding is appropriately pre-trained with unlabeled data, its inclusion is a form of semi-supervised learning and could be useful. We will show later, however, that this type of approach falls behind our approach of learning *region embeddings* through training one-hot LSTM on unlabeled data. Altogether, elimination of the word embedding layer was found to be useful; thus, we base our work on one-hot LSTM.

2.2. More simplifications

We introduce four more useful modifications to wv-LSTM that lead to higher accuracy or faster training.

Pooling: simplifying sub-problems Our framework of ‘region embedding + pooling’ has a simplification effect as follows. In wv-LSTM, the sub-problem that LSTM needs to solve is to represent the entire document by one vector (*document embedding*). We make this easy by changing it to detecting regions of text (of arbitrary sizes) that are relevant to the task and representing them by vectors (*region embedding*). As illustrated in Figure 3, we let the LSTM layer emit vectors \mathbf{h}_t at each time step, and let pooling

aggregate them into a document vector. With *wv-LSTM*, LSTM has to remember relevant information until it gets to the end of the document even if relevant information was observed 10K words away. The task of our LSTM is easier as it is allowed to forget old things via the forget gate and can focus on representing the concepts conveyed by smaller segments such as phrases or sentences.

A related architecture appears in the Deep Learning Tutorials² though it uses a word embedding. Another related work is (Lai et al., 2015), which combined pooling with non-LSTM recurrent networks and a word embedding.

Chopping for speeding up training In addition to simplifying the sub-problem, pooling has the merit of enabling faster training via *chopping*. Since we set the goal of LSTM to embedding text regions instead of documents, it is no longer crucial to go through the document from the beginning to the end sequentially. At the time of training, we can *chop* each document into segments of a fixed length that is sufficiently long (e.g., 50 or 100) and process all the segments in a mini batch in parallel as if these segments were individual documents. (Note that this is done only in the LSTM layer and pooling is done over the entire document.) We perform testing without chopping. That is, we train LSTM with approximations of sequences for speed up and test with real sequences for better accuracy. There is a risk of chopping important phrases (e.g., “don’t | like it”), and this can be easily avoided by having segments slightly overlap. However, we found that gains from overlapping segments tend to be small and so our experiments reported below were done without overlapping.

Removing the input/output gates We found that when LSTM is followed by pooling, the presence of input and output gates typically does not improve accuracy, while removing them nearly halves the time and memory required for training and testing. It is intuitive, in particular, that pooling can make the output gate unnecessary; the role of the output gate is to prevent undesirable information from entering the output \mathbf{h}_t , and such irrelevant information can be filtered out by max-pooling. Without the input and output gates, the LSTM formulation can be simplified to:

$$\mathbf{f}_t = \sigma(\mathbf{W}^{(f)}\mathbf{x}_t + \mathbf{U}^{(f)}\mathbf{h}_{t-1} + \mathbf{b}^{(f)}), \quad (2)$$

$$\mathbf{u}_t = \tanh(\mathbf{W}^{(u)}\mathbf{x}_t + \mathbf{U}^{(u)}\mathbf{h}_{t-1} + \mathbf{b}^{(u)}), \quad (3)$$

$$\mathbf{c}_t = \mathbf{u}_t + \mathbf{f}_t \odot \mathbf{c}_{t-1}, \quad \mathbf{h}_t = \tanh(\mathbf{c}_t).$$

This is equivalent to fixing \mathbf{i}_t and \mathbf{o}_t to all ones. It is in spirit similar to Gated Recurrent Units (Cho et al., 2014) but simpler, having fewer gates.

Bidirectional LSTM for better accuracy The changes from *wv-LSTM* above substantially reduce the time and

	Chop	Time	Error
1-layer oh-CNN	–	18	7.64
wv-LSTM	–	337	11.59
wv-LSTMp	100	110	10.90
oh-LSTMp	100	88	7.72
oh-LSTMp; no i/o gates	100	48	7.68
oh-2LSTMp; no i/o gates	50	84	7.33

Table 1. Training time and error rates of LSTMs on Elec. “Chop”: chopping size. “Time”: seconds per epoch for training on Tesla M2070. “Error”: classification error rates (%) on test data. “wv-LSTMp”: word-vector LSTM with pooling. “oh-LSTMp”: one-hot LSTM with pooling. “oh-2LSTMp”: one-hot bidirectional LSTM with pooling.

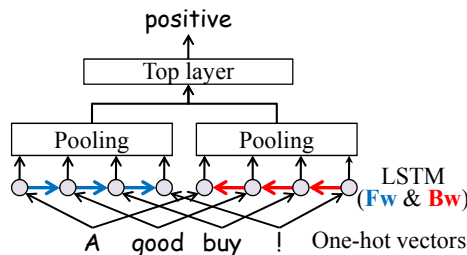


Figure 4. *oh-2LSTMp*: our one-hot bidirectional LSTM with pooling.

memory required for training and make it practical to add one more layer of LSTM going in the opposite direction for accuracy improvement. As shown in Figure 4, we concatenate the output of a forward LSTM (left to right) and a backward LSTM (right to left), which is referred to as *bidirectional LSTM* in the literature. The resulting model is a *one-hot bidirectional LSTM with pooling*, and we abbreviate it to *oh-2LSTMp*. Table 1 shows how much accuracy and/or training speed can be improved by elimination of the word embedding layer, pooling, chopping, removing the input/output gates, and adding the backward LSTM.

2.3. Experiments (supervised)

We used four datasets: IMDB, Elec, RCV1 (second-level topics), and 20-newsgroup (20NG)³, to facilitate direct comparison with JZ15 and DL15. The first three were used in JZ15. IMDB and 20NG were used in DL15. The datasets are summarized in Table 2.

The data was converted to lower-case letters. In the neural network experiments, vocabulary was reduced to the most frequent 30K words of the training data to reduce computational burden; square loss was minimized with dropout (Hinton et al., 2012) applied to the input to the top layer; weights were initialized by the Gaussian distribution with

² <http://deeplearning.net/tutorial/lstm.html>

³ <http://ana.cachopo.org/datasets-for-single-label-text-categorization>

	#train	#test	avg	max	#class
IMDB	25,000	25,000	265	3K	2
Elec	25,000	25,000	124	6K	2
RCV1	15,564	49,838	249	12K	55
20NG	11,293	7,528	267	12K	20

Table 2. Data. “avg”/“max”: the average/maximum length of documents (#words) of the training/test data. IMDB and Elec are for sentiment classification (positive vs. negative) of movie reviews and Amazon electronics product reviews, respectively. RCV1 (second-level topics only) and 20NG are for topic categorization of Reuters news articles and newsgroup messages, respectively.

zero mean and standard deviation 0.01. Optimization was done with SGD with mini-batch size 50 or 100 with momentum or optionally *rmsprop* (Tieleman & Hinton, 2012) for acceleration.

Hyper parameters such as learning rates were chosen based on the performance on the development data, which was a held-out portion of the training data, and training was re-done using all the training data with the chosen parameters.

We used the same pooling method as used in JZ15, which parameterizes the number of pooling regions so that pooling is done for k non-overlapping regions of equal size, and the resulting k vectors are concatenated to make one vector per document. The pooling settings chosen based on the performance on the development data are the same as JZ15a, which are max-pooling with $k=1$ on IMDB and Elec and average-pooling with $k=10$ on RCV1; on 20NG, max-pooling with $k=10$ was chosen.

methods	IMDB	Elec	RCV1	20NG
SVM bow	11.36	11.71	10.76	17.47
SVM 1–3grams	9.42	8.71	10.69	15.85
wv-LSTM [DL15]	13.50	11.74	16.04	18.0
oh-2LSTMp	8.14	7.33	11.17	13.32
oh-CNN [JZ15b]	8.39	7.64	9.17	13.64

Table 3. Error rates (%). Supervised results without any pre-training. SVM and oh-CNN results on all but 20NG are from JZ15a and JZ15b, respectively; wv-LSTM results on IMDB and 20NG are from DL15; all others are new experimental results of this work.

Table 3 shows the error rates obtained without any additional unlabeled data or pre-training of any sort. For meaningful comparison, this table shows neural networks with comparable dimensionality of embeddings, which are one-hot CNN with one convolution layer with 1000 feature maps and bidirectional LSTMs of 500 units each. In other words, the convolution layer produces a 1000-dimensional vector at each location, and the LSTM in each direction

emits a 500-dimensional vector at each time step. An exception is wv-LSTM, equipped with 512 LSTM units (smaller than 2×500) and a word embedding layer of 512 dimensions; DL15 states that without pre-training, addition of more LSTM units broke down training. A more complex and larger one-hot CNN will be reviewed later.

Comparing the two types of LSTM in Table 3, we see that our one-hot bidirectional LSTM with pooling (oh-2LSTMp) outperforms word-vector LSTM (wv-LSTM) on all the datasets, confirming the effectiveness of our approach.

Now we review the non-LSTM baseline methods. The last row of Table 3 shows the best one-hot CNN results within the constraints above. They were obtained by *bow-CNN* (whose input to the embedding function (1) is a bow vector of the region) with region size 20 on RCV1, and *seq-CNN* (with the regular concatenation input) with region size 3 on the others. In Table 3, on three out of the four datasets, oh-2LSTMp outperforms SVM and the CNN. However, on RCV1, it underperforms both. We conjecture that this is because strict word order is not very useful on RCV1. This point can also be observed in the SVM and CNN performances. Only on RCV1, n -gram SVM is no better than bag-of-words SVM, and only on RCV1, bow-CNN outperforms seq-CNN. That is, on RCV1, bags of words in a window of 20 at every location are more useful than words in strict order. This is presumably because the former can more easily cover variability of expressions indicative of topics. Thus, LSTM, which does not have an ability to put words into bags, loses to bow-CNN.

methods	IMDB	Elec	20NG
oh-2LSTMp, copied from Tab.3	8.14	7.33	13.32
oh-CNN, 2 region sizes [JZ15a]	8.04	7.48	13.55

More on one-hot CNN vs. one-hot LSTM LSTM can embed regions of variable (and possibly large) sizes whereas CNN requires the region size to be fixed. We attribute to this fact the small improvements of oh-2LSTMp over oh-CNN in Table 3. However, this shortcoming of CNN can be alleviated by having multiple convolution layers with distinct region sizes. We show in the table above that one-hot CNNs with two layers (of 1000 feature maps each) with two different region sizes⁴ rival oh-2LSTMp. Although these models are larger than those in Table 3, training/testing is still faster than the LSTM models due to simplicity of the region embeddings. By comparison, the strength of LSTM to embed larger regions appears not to be a big contributor here. This may be because the amount of training data is not sufficient enough to learn the relevance of longer word sequences. Overall, one-hot CNN works

⁴ Region sizes were 2 and 3 for IMDB, 3 and 4 for Elec, and 3 and 20 (bow input) for 20NG.

surprising well considering its simplicity, and this observation motivates the idea of *combining the two types of region embeddings*, discussed later.

Comparison with the previous best results on 20NG

The previous best performance on 20NG is 15.3 (not shown in the table) of DL15, obtained by *pre-training* *wv-LSTM* of 1024 units with labeled training data. Our *oh-2LSTMp* achieved 13.32, which is 2% better. The previous best results on the other datasets use unlabeled data, and we will review them with our semi-supervised results.

3. Semi-supervised LSTM

To exploit unlabeled data as an additional resource, we use a non-linear extension of *two-view feature learning*, whose linear version appeared in our earlier work (Ando & Zhang, 2005; 2007). This was used in JZ15b to learn from unlabeled data a region embedding embodied by a convolution layer. In this work we use it to learn a region embedding embodied by a one-hot LSTM. Let us start with a brief review of non-linear two-view feature learning.

3.1. Two-view embedding (tv-embedding) [JZ15b]

A rough sketch is as follows. Consider two views of the input. An embedding is called a *tv-embedding* if the embedded view is as good as the original view for the purpose of predicting the other view. If the two views and the labels (classification targets) are related to one another only through some hidden states, then the tv-embedded view is as good as the original view for the purpose of classification. Such an embedding is useful provided that its dimensionality is much lower than the original view.

JZ15b applied this idea by regarding text regions embedded by the convolution layer as one view and their surrounding context as the other view and training a tv-embedding (embodied by a convolution layer) on unlabeled data. The obtained tv-embeddings were used to produce additional input to a supervised region embedding of one-hot CNN, resulting in higher accuracy.

3.2. Learning LSTM tv-embeddings

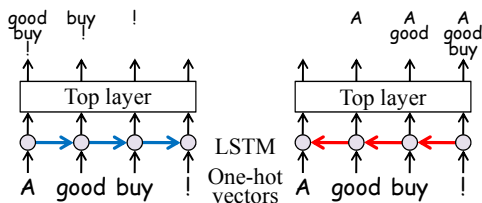


Figure 5. Training LSTM tv-embeddings on unlabeled data

In this work we obtain a tv-embedding in the form of LSTM from unlabeled data as follows. At each time step,

IMDB	75K	(20M words)	Provided
Elec	200K	(24M words)	Provided
RCV1	669K	(183M words)	Sept'96-June'97

Table 4. Unlabeled data. See JZ15b for more details.

we consider the following two views: the words we have already seen in the document (view-1), and the next few words (view-2). The task of tv-embedding learning is to predict view-2 based on view-1. We train one-hot LSTMs in both directions, as in Figure 5, on unlabeled data. For this purpose, we use the input and output gates as well as the forget gate as we found them to be useful.

The theory of tv-embedding says that the region embeddings obtained in this way are useful for the task of interest *if* the two views are related to each other through the concepts relevant to the task. To reduce undesirable relations between the views such as syntactic relations, JZ15b performed vocabulary control to remove function words from (and only from) the vocabulary of the target view, which we found useful also for LSTM.

We use the tv-embeddings obtained from unlabeled data to produce *additional input* to LSTM by replacing (2) and (3) by the following:

$$\mathbf{f}_t = \sigma(\mathbf{W}^{(f)} \mathbf{x}_t + \sum_{j \in S} \widetilde{\mathbf{W}}^{(j,f)} \tilde{\mathbf{x}}_t^j + \mathbf{U}^{(f)} \mathbf{h}_{t-1} + \mathbf{b}^{(f)}),$$

$$\mathbf{u}_t = \sigma(\mathbf{W}^{(u)} \mathbf{x}_t + \sum_{j \in S} \widetilde{\mathbf{W}}^{(j,u)} \tilde{\mathbf{x}}_t^j + \mathbf{U}^{(u)} \mathbf{h}_{t-1} + \mathbf{b}^{(u)}).$$

$\tilde{\mathbf{x}}_t^j$ is the output of a tv-embedding (an LSTM trained with unlabeled data) indexed by j at time step t , and S is a set of tv-embeddings which contains the two LSTMs going forward and backward as in Figure 5. Although it is possible to fine-tune the tv-embeddings with labeled data, for simplicity and faster training, we fixed them in our experiments.

3.3. Combining LSTM tv-embeddings and CNN tv-embeddings

It is easy to see that the set S above can be expanded with *any tv-embeddings*, not only those in the form of LSTM (*LSTM tv-embeddings*) but also with the tv-embeddings in the form of convolution layers (*CNN tv-embeddings*) such as those obtained in JZ15b. Similarly, it is possible to use LSTM tv-embeddings to produce additional input to CNN. While both LSTM tv-embeddings and CNN tv-embeddings are region embeddings, their formulations are very different from each other; therefore, we expect that they complement each other and bring further performance improvements when combined. We will empirically confirm these conjectures in the experiments below. Note that being able to naturally combine several tv-embeddings is a strength of

		Unlabeled data usage	IMDB	Elec	RCV1
1	wv-LSTM [DL15]	Pre-training	7.24	6.84	14.65
2	wv-2LSTMp	300-dim Google News word2vec	8.67	7.64	10.62
3		200-dim word2vec scaled	7.29	6.76	10.18
4	oh-2LSTMp	2×100-dim LSTM tv-embed.	6.66	6.08	9.24
5	oh-CNN [JZ15b]	1×200-dim CNN tv-embed.	6.81	6.57	7.97

Table 5. Semi-supervised error rates (%). The wv-LSTM result on IMDB is from [DL15]; the oh-CNN results are from [JZ15b]; all others are the results of our new experiments.

our framework, which uses unlabeled data to produce *additional input* to LSTM instead of pre-training.

3.4. Semi-supervised experiments

We used IMDB, Elec, and RCV1 for our semi-supervised experiments; 20NG was excluded due to the absence of standard unlabeled data. Table 4 summarizes the unlabeled data. To experiment with LSTM tv-embeddings, we trained two LSTMs (forward and backward) with 100 units each on unlabeled data. The training objective was to predict the next k words where k was set to 20 for RCV1 and 5 for others. Similar to JZ15b, we minimized weighted square loss $\sum_{i,j} \alpha_{i,j} (\mathbf{z}_i[j] - \mathbf{p}_i[j])^2$ where i goes through the time steps, \mathbf{z} represents the next k words by a bow vector, and \mathbf{p} is the model output; $\alpha_{i,j}$ were set to achieve negative sampling effect for speed-up; vocabulary control was performed for reducing undesirable relations between views, which sets the vocabulary of the target (i.e., the k words) to the 30K most frequent words excluding function words (or stop words on RCV1). Other details followed the supervised experiments.

Our semi-supervised one-hot bidirectional LSTM with pooling (oh-2LSTMp) in row#4 of Table 5 used the two LSTM tv-embeddings trained on unlabeled data as described above, to produce additional input to one-hot LSTMs in two directions (500 units each). Compared with the supervised oh-2LSTMp (Table 3), clear performance improvements were obtained on all the datasets, thus, confirming the effectiveness of our approach.

We review the semi-supervised performance of wv-LSTMs (Table 5 row#1). In DL15 the model consisted of a word embedding layer of 512 dimensions, an LSTM layer with 1024 units, and 30 hidden units on top of the LSTM layer; the word embedding layer and the LSTM were pre-trained with unlabeled data and were fine-tuned with labeled data; pre-training used either the language model objective or auto-encoder objective. The error rate on IMDB is from DL15, and those on Elec and RCV1 are our best effort to perform pre-training with the language model objective. We used the same configuration on Elec as DL15; however, on RCV1, which has 55 classes, 30 hidden units turned out to be too few and we changed it to 1000. Although the pre-

trained wv-LSTM clearly outperformed the supervised wv-LSTM (Table 3), it underperformed the models with region tv-embeddings (Table 5 row#4,5).

Previous studies on LSTM for text often convert words into pre-trained word vectors, and word2vec (Mikolov et al., 2013) is a popular choice for this purpose. Therefore, we tested wv-2LSTMp (word-vector bidirectional LSTM with pooling), whose only difference from oh-2LSTMp is that the input to the LSTM layers is the pre-trained word vectors. The word vectors were optionally updated (fine-tuned) during training. Two types of word vectors were tested. The Google News word vectors were trained by word2vec on a huge (100 billion-word) news corpus and are provided publicly. On our tasks, wv-2LSTMp using the Google News vectors (Table 5 row#2) performed relatively poorly. When word2vec was trained with the domain unlabeled data, better results were observed after we scaled word vectors appropriately (Table 5 row#3). Still, it underperformed the models with region tv-embeddings (row #4,5), which used the same domain unlabeled data. We attribute the superiority of the models with tv-embeddings to the fact that they learn, from unlabeled data, embeddings of *text regions*, which can convey higher-level concepts than single words in isolation.

Now we review the performance of one-hot CNN with one 200-dim CNN tv-embedding (Table 5 row#5), which is comparable with our LSTM with two 100-dim LSTM tv-embeddings (row#4) in terms of the dimensionality of tv-embeddings. The LSTM (row#4) rivals or outperforms the CNN (row#5) on IMDB/Elec but underperforms it on RCV1. Increasing the dimensionality of LSTM tv-embeddings from 100 to 300 on RCV1, we obtain 8.62, but it still does not reach 7.97 of the CNN. As discussed earlier, we attribute the superiority of one-hot CNN on RCV1 to its unique way of representing parts of documents via bow input.

3.5. Experiments combining CNN tv-embeddings and LSTM tv-embeddings

In Section 3.3 we noted that LSTM tv-embeddings and CNN tv-embeddings can be naturally combined. We experimented with this idea in the following two settings.

		Unlabeled data usage	IMDB	Elec	RCV1
1	oh-2LSTMp	two LSTM tv-embed.	6.66	6.08	8.62
2	oh-CNN [JZ15b]	3×100-dim CNN tv-embed.	6.51	6.27	7.71
3	oh-2LSTMp	3×100-dim CNN tv-embed.	5.94	5.55	8.52
4	oh-CNN	+ two LSTM tv-embed.	6.05	5.87	7.15

Table 6. Error rates (%) obtained by combining CNN tv-embed. and LSTM tv-embed. (rows 3–4). LSTM tv-embed. were 100-dim each on IMDB and Elec, and 300-dim on RCV1. To see the combination effects, compare row#3 with #1, and compare row#4 with #2.

	U	IMDB	Elec	RCV1
oh-CNN+doc. [JZ15a]	N	7.67	7.14	–
Co-tr. <i>optimized</i> [JZ15b]	Y	(8.06)	(7.63)	(8.73)
Para.vector [LM14]	Y	7.42	–	–
wv-LSTM [DL15]	Y	7.24	–	–
oh-CNN(semi.) [JZ15b]	Y	6.51	6.27	7.71
Our best model	Y	5.94	5.55	7.15

Table 7. Comparison with previous best results. Error rates (%). “U”: Was unlabeled data used? “Co-tr. optimized”: co-training using oh-CNN as a base learner with parameters (e.g., when to stop) *optimized on the test data*; it demonstrates the difficulty of exploiting unlabeled data on these tasks.

In one setting, oh-2LSTMp takes additional input from five embeddings: two LSTM tv-embeddings used in Table 5 and three CNN tv-embeddings from JZ15b obtained by three distinct combinations of training objectives and input representations, which are publicly provided. These CNN tv-embeddings were trained to be applied to text regions of size k at every location taking bow input, where k is 5 on IMDB/Elec and 20 on RCV1. We connect each of the CNN tv-embeddings to an LSTM by aligning the centers of the regions of the former with the LSTM time steps; e.g., the CNN tv-embedding result on the first five words is passed to the LSTM at the time step on the third word. In the second setting, we trained one-hot CNN with these five types of tv-embeddings by replacing (1) $\max(0, \mathbf{W}\mathbf{x}_\ell + \mathbf{b})$ by $\max(0, \mathbf{W}\mathbf{x}_\ell + \sum_j \widetilde{\mathbf{W}}^{(j)}\tilde{\mathbf{x}}_\ell^j + \mathbf{b})$ where $\tilde{\mathbf{x}}_\ell^j$ is the output of the j -th tv-embedding with the same alignment as above.

Rows 3–4 of Table 6 show the results of these two types of models. For comparison, we also show the results of the LSTM with LSTM tv-embeddings only (row#1) and the CNN with CNN tv-embeddings only (row#2). To see the effects of combination, compare row#3 with row#1, and compare row#4 with row#2. For example, adding the CNN tv-embeddings to the LSTM of row#1, the error rate on IMDB improved from 6.66 to 5.94, and adding the LSTM tv-embeddings to the CNN of row#2, the error rate on RCV1 improved from 7.71 to 7.15. The results indicate that, as expected, LSTM tv-embeddings and CNN tv-embeddings complement each other and improve performance when combined.

3.6. Comparison with the previous best results

The previous best results in the literature are shown in Table 7. More results of previous semi-supervised models can be found in JZ15b, all of which clearly underperform the semi-supervised one-hot CNN of Table 7. The best supervised results on IMDB/Elec of JZ15a are in the first row, obtained by integrating a document embedding layer into one-hot CNN. Many more of the previous results on IMDB can be found in (Le & Mikolov, 2014), all of which are over 10% except for 8.78 by bi-gram NBSVM (Wang & Manning, 2012). 7.42 by paragraph vectors (Le & Mikolov, 2014) and 6.51 by JZ15b were considered to be large improvements. As shown in the last row of Table 7, our new model further improved it to 5.94; also on Elec and RCV1, our best models exceeded the previous best results.

4. Conclusion

Within the general framework of ‘region embedding + pooling’ for text categorization, we explored region embeddings via one-hot LSTM. The region embedding of one-hot LSTM rivaled or outperformed that of the state-of-the-art one-hot CNN, proving its effectiveness. We also found that the models with either one of these two types of region embedding strongly outperformed other methods including previous LSTM. The best results were obtained by combining the two types of region embedding trained on unlabeled data, suggesting that their strengths are complementary. As a result, we reported substantial improvements over the previous best results on benchmark datasets.

At a high level, our results indicate the following. First, on this task, embeddings of text regions, which can convey higher-level concepts, are more useful than embeddings of single words in isolation. Second, useful region embeddings can be learned by working with one-hot vectors directly, either on labeled data or unlabeled data. Finally, a promising future direction might be to seek, under this framework, new region embedding methods with complementary benefits.

Acknowledgements

We would like to thank anonymous reviewers for valuable feedback. This research was partially supported by NSF IIS-1250985, NSF IIS-1407939, and NIH R01AI116744.

References

- Ando, Rie K. and Zhang, Tong. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6:1817–1853, 2005.
- Ando, Rie K. and Zhang, Tong. Two-view feature generation model for semi-supervised learning. In *Proceedings of ICML, 2007*.
- Cho, Kyunghyun, van Merriënboer, Bart, Gulcehre, Caglar, Bahdanau, Dzmitry, Bougares, Fethi, Schwenk, Holger, and Bengio, Yoshua. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of EMNLP, 2014*.
- Dai, Andrew M. and Le, Quoc V. Semi-supervised sequence learning. In *NIPS, 2015*.
- Gers, Felix A., Schmidhuder, Jürgen, and Cummins, Fred. Learning to forget: Continual prediction with LSTM. *Neural Computation*, 12(10):2451–2471, 2000.
- Hinton, Geoffrey E., Srivastava, Nitish, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan R. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv:1207.0580*, 2012.
- Hochreiter, Sepp and Schmidhuder, Jürgen. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- Joachims, Thorsten. Text categorization with support vector machines: Learning with many relevant features. In *ECML, 1998*.
- Johnson, Rie and Zhang, Tong. Effective use of word order for text categorization with convolutional neural networks. In *NAACL HLT, 2015a*.
- Johnson, Rie and Zhang, Tong. Semi-supervised convolutional neural networks for text categorization via region embedding. In *NIPS, 2015b*.
- Lai, Siwei, Xu, Liheng, Liu, Kang, and Zhao, Jun. Recurrent convolutional neural networks for text classification. In *Proceedings of AAAI, 2015*.
- Le, Phong and Zuidema, Willem. Compositional distributional semantics with long short-term memory. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics, 2015*.
- Le, Quoc and Mikolov, Tomas. Distributed representations of sentences and documents. In *Proceedings of ICML, 2014*.
- LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE, 1986*.
- Lewis, David D., Yang, Yiming, Rose, Tony G., and Li, Fan. RCV1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, 5:361–397, 2004.
- Mikolov, Tomas, Sutskever, Ilya, Chen, Kai, Corrado, Greg, and Dean, Jeffrey. Distributed representations of words and phrases and their compositionality. In *NIPS, 2013*.
- Sutskever, Hya, Vinyals, Oriol, and Le, Quoc V. Sequence to sequence learning with neural networks. In *NIPS, 2014*.
- Tai, Kai Sheng, Socher, Richard, and Manning, Christopher D. Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of ACL, 2015*.
- Tang, Duyu, Qin, Bing, and Liu, Ting. Document modeling with gated recurrent neural network for sentiment classification. In *Proceedings of EMNLP, 2015*.
- Tieleman, Tijman and Hinton, Geoffrey. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4, 2012.
- Wang, Sida and Manning, Christopher D. Baselines and bigrams: Simple, good sentiment and topic classification. In *Proceedings of ACL*, pp. 90–94, 2012.
- Zaremba, Wojciech and Sutskever, Ilya. Learning to execute. *arXiv:1410.4615*, 2014.
- Zhang, Xiang, Zhao, Junbo, and LeCun, Yann. Character-level convolutional networks for text classification. In *NIPS, 2015*.
- Zhu, Xiaodan, Sobhani, Parinaz, and Guo, Hongyu. Long short-term memory over recursive structures. In *Proceedings of ICML, 2015*.