

Speeding up Context-based Sentence Representation Learning with Non-autoregressive Convolutional Decoding

Shuai Tang¹, Hailin Jin², Chen Fang², Zhaowen Wang², Virginia R. de Sa¹

¹ Department of Cognitive Science, UC San Diego

² Adobe Research

{shuaitang93, desa}@ucsd.edu, {hljin, cfang, zhawang}@adobe.com

Abstract

Context plays an important role in human language understanding, thus it may also be useful for machines learning vector representations of language. In this paper, we explore an asymmetric encoder-decoder structure for unsupervised context-based sentence representation learning. We carefully designed experiments to show that neither an autoregressive decoder nor an RNN decoder is required. After that, we designed a model which still keeps an RNN as the encoder, while using a non-autoregressive convolutional decoder. We further combine a suite of effective designs to significantly improve model efficiency while also achieving better performance. Our model is trained on two different large unlabelled corpora, and in both cases the transferability is evaluated on a set of downstream NLP tasks. We empirically show that our model is simple and fast while producing rich sentence representations that excel in downstream tasks.

1 Introduction

Learning distributed representations of sentences is an important and hard topic in both the deep learning and natural language processing communities, since it requires machines to encode a sentence with rich language content into a fixed-dimension vector filled with real numbers. Our goal is to build a distributed sentence encoder learnt in an unsupervised fashion by exploiting the structure and relationships in a large unlabelled corpus.

Numerous studies in human language processing have supported that rich semantics of a word or sentence can be inferred from its context (Harris, 1954; Firth, 1957). The idea of learning from the

co-occurrence (Turney and Pantel, 2010) was recently successfully applied to vector representation learning for words in Mikolov et al. (2013) and Pennington et al. (2014).

A very recent successful application of the distributional hypothesis (Harris, 1954) at the sentence-level is the skip-thoughts model (Kiros et al., 2015). The skip-thoughts model learns to encode the current sentence and decode the surrounding two sentences instead of the input sentence itself, which achieves overall good performance on all tested downstream NLP tasks that cover various topics. The major issue is that the training takes too long since there are two RNN decoders to reconstruct the previous sentence and the next one independently. Intuitively, given the current sentence, inferring the previous sentence and inferring the next one should be different, which supports the usage of two independent decoders in the skip-thoughts model. However, Tang et al. (2017) proposed the skip-thought neighbour model, which only decodes the next sentence based on the current one, and has similar performance on downstream tasks compared to that of their implementation of the skip-thoughts model.

In the encoder-decoder models for learning sentence representations, only the encoder will be used to map sentences to vectors after training, which implies that the quality of the generated language is not our main concern. This leads to our two-step experiment to check the necessity of applying an autoregressive model as the decoder. In other words, since the decoder’s performance on language modelling is not our main concern, it is preferred to reduce the complexity of the decoder to speed up the training process. In our experiments, the first step is to check whether “teacher-forcing” is required during training if we stick to using an autoregressive model as the decoder, and the second step is to check whether an autoregres-

sive decoder is necessary to learn a good sentence encoder. Briefly, the experimental results show that an autoregressive decoder is indeed not essential in learning a good sentence encoder; thus the two findings of our experiments lead to our final model design.

Our proposed model has an asymmetric encoder-decoder structure, which keeps an RNN as the encoder and has a CNN as the decoder, and the model explores using only the subsequent context information as the supervision. The asymmetry in both model architecture and training pair reduces a large amount of the training time.

The contribution of our work is summarised as:

1. We design experiments to show that an autoregressive decoder or an RNN decoder is not necessary in the encoder-decoder type of models for learning sentence representations, and based on our results, we present two findings. **Finding I:** It is not necessary to input the correct words into an autoregressive decoder for learning sentence representations. **Finding II:** The model with an autoregressive decoder performs similarly to the model with a predict-all-words decoder.
2. The two findings above lead to our final model design, which keeps an RNN encoder while using a CNN decoder and learns to encode the current sentence and decode the subsequent contiguous words all at once.
3. With a suite of techniques, our model performs decently on the downstream tasks, and can be trained efficiently on a large unlabelled corpus.

The following sections will introduce the components in our ‘‘RNN-CNN’’ model, and discuss our experimental design.

2 RNN-CNN Model

Our model is highly asymmetric in terms of both the training pairs and the model structure. Specifically, our model has an RNN as the encoder, and a CNN as the decoder. During training, the encoder takes the i -th sentence S_i as the input, and then produces a fixed-dimension vector \mathbf{z}_i as the sentence representation; the decoder is applied to reconstruct the paired target sequence T_i that contains the subsequent contiguous words. The distance between

the generated sequence and the target one is measured by the cross-entropy loss at each position in T_i . An illustration is in Figure 1. (For simplicity, we omit the subscript i in this section.)

1. Encoder: The encoder is a bi-directional Gated Recurrent Unit (GRU, Chung et al. (2014))¹. Suppose that an input sentence S contains M words, which are $\{w_s^1, w_s^2, \dots, w_s^M\}$, and they are transformed by an embedding matrix $\mathbf{E} \in \mathbb{R}^{d_e \times |V|}$ to word vectors². The bi-directional GRU takes one word vector at a time, and processes the input sentence in both the forward and backward directions; both sets of hidden states are concatenated to form the hidden state matrix $\mathbf{H} = [\mathbf{h}^1, \mathbf{h}^2, \dots, \mathbf{h}^M] \in \mathbb{R}^{d_h \times M}$, where d_h is the dimension of the hidden states $\mathbf{h}^m = [\overleftarrow{\mathbf{h}}^m; \overrightarrow{\mathbf{h}}^m]$ ($\forall m \in \{1, 2, \dots, M\}$).

2. Representation: We aim to provide a model with faster training speed and better transferability than existing algorithms; thus we choose to apply a parameter-free composition function, which is a concatenation of the outputs from a global mean pooling over time and a global max pooling over time, on the computed sequence of hidden states \mathbf{H} . The composition function is represented as

$$\mathbf{z} = \left[\frac{1}{M} \sum_{m=1}^M \mathbf{h}^m; \text{MaxPool}(\mathbf{H}) \right] \quad (1)$$

where MaxPool is the max operation on each row of the matrix \mathbf{H} , which outputs a vector with dimension d_h . Thus the representation $\mathbf{z} \in \mathbb{R}^{2d_h}$.

3. Decoder: The decoder is a 3-layer CNN to reconstruct the paired target sequence T , which needs to expand \mathbf{z} , which can be considered as a sequence with only one element, to a sequence with T elements. Intuitively, the decoder could be a stack of deconvolution layers. For fast training speed, we optimised the architecture to make it possible to use fully-connected layers and convolution layers in the decoder, since generally, convolution layers run faster than deconvolution layers in modern deep learning frameworks.

Suppose that the target sequence T has N words, which are $\{w_t^1, w_t^2, \dots, w_t^N\}$, the first layer of deconvolution will expand $\mathbf{z} \in \mathbb{R}^{2d_h \times 1}$, into a feature

¹We experimented with both Long-short Term Memory (LSTM, Hochreiter and Schmidhuber (1997)) and GRU. Since LSTM didn’t give us significant performance boost, and generally GRU runs faster than LSTM, in our experiments, we stick to using GRU in the encoder.

² V is the vocabulary, and $|V|$ is the number of unique words in the vocabulary. d_e is the dimension of each word vector.




Figure 1: Our proposed model is composed of an RNN encoder, and a CNN decoder. During training, a batch of sentences are sent to the model, and the RNN encoder computes a vector representation for each of sentences; then the CNN decoder needs to reconstruct the paired target sequence, which contains 30 contiguous words right after the input sentence, given the vector representation. 300 is the dimension of word vectors. $2d_h$ is the dimension of the sentence representation which varies with the RNN encoder size. (Better view in colour.)

map with N elements. It can be easily implemented as a concatenation of outputs from N linear transformations in parallel. Then the second and third layer are 1D-convolution layers. The output feature map is $\mathbf{U} = [\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^N] \in \mathbb{R}^{d_e \times N}$, where d_e is the dimension of the word vectors.

Note that our decoder is not an autoregressive model and has high training efficiency. We will discuss the reason for choosing this decoder which we call a predict-all-words CNN decoder.

4. Objective: The training objective is to maximise the likelihood of the target sequence being generated from the decoder. Since in our model, each word is predicted independently, a softmax layer is applied after the decoder to produce a probability distribution over words in V at each position, thus the probability of generating a word w_t^n in the target sequence is defined as:

$$P(w_t^n) = \frac{e^{\mathbf{E}(w_t^n)^\top \mathbf{u}^n}}{\sum_{w \in V} e^{\mathbf{E}(w)^\top \mathbf{u}^n}}, \quad (2)$$

where, $\mathbf{E}(w)$ is the vector representation of w in the embedding matrix \mathbf{E} , and $\mathbf{E}(w)^\top \mathbf{u}^n$ is the dot-product between the word vector and the feature vector produced by the decoder at position n . The training objective is to minimise the sum of the negative log-likelihood over all positions in the target sequence T :

$$\begin{aligned} \mathcal{L}(\phi, \theta) &= -\log P(T|S; \phi, \theta) \\ &= -\sum_{n=1}^N \log P(w_t^n | w_s^1, w_s^2, \dots, w_s^M; \phi, \theta), \quad (3) \end{aligned}$$

where ϕ and θ contain the parameters in the encoder and the decoder, respectively. The training objective $\mathcal{L}(\phi, \theta)$ is summed over all sentences in the training corpus.

3 Architecture Design

We use an encoder-decoder model and use context for learning sentence representations in an unsupervised fashion. Since the decoder won't be used after training, and the quality of the generated sequences is not our main focus, it is important to study the design of the decoder. Generally, a fast training algorithm is preferred; thus proposing a new decoder with high training efficiency and also strong transferability is crucial for an encoder-decoder model.

3.1 CNN as the decoder

Our design of the decoder is basically a 3-layer ConvNet that predicts all words in the target sequence at once. In contrast, existing work, such as skip-thoughts (Kiros et al., 2015), and CNN-LSTM (Gan et al., 2017), use autoregressive RNNs as the decoders. As known, an autoregressive model is good at generating sequences with high quality, such as language and speech. However, an autoregressive decoder seems to be unnecessary in an encoder-decoder model for learning sentence representations, since it won't be used after training, and it takes up a large portion of the training time to compute the output and the gradient. Therefore, we conducted experiments to test the necessity of using an autoregressive decoder in learning sentence representations, and we had two findings.

Decoder	SICK-R	SICK-E	STS14	MSRP (Acc/F1)	SST	TREC
auto-regressive RNN as decoder						
Teacher-Forcing	0.8530	82.6	0.51/0.50	74.1 / 81.7	82.5	88.2
Always Sampling	0.8576	83.2	0.55/0.53	74.7 / 81.3	80.6	87.0
Uniform Sampling	0.8559	82.9	0.54/0.53	74.0 / 81.8	81.0	87.4
auto-regressive CNN as decoder						
Teacher-Forcing	0.8510	82.8	0.49/0.48	74.7 / 82.8	81.4	82.6
Always Sampling	0.8535	83.3	0.53/0.52	75.0 / 81.7	81.4	87.6
Uniform Sampling	0.8568	83.4	0.56/0.54	74.7 / 81.4	83.0	88.4
predict-all-words RNN as decoder						
RNN	0.8508	82.8	0.58/0.55	74.2 / 82.8	81.6	88.8
predict-all-words CNN as decoder						
CNN	0.8530	82.6	0.58/0.56	75.6 / 82.9	82.8	89.2
CNN-MaxOnly	0.8465	82.6	0.50/0.47	73.3 / 81.5	79.1	82.2
Double-sized RNN Encoder						
CNN	0.8631	83.9	0.58/0.55	74.7 / 83.1	83.4	90.2
CNN-MaxOnly	0.8485	83.2	0.47/0.44	72.9 / 80.8	82.2	86.6

Table 1: The models here all have a bi-directional GRU as the encoder (dimensionality 300 in each direction). The default way of producing the representation is a concatenation of outputs from a global *mean-pooling* and a global *max-pooling*, while “*-MaxOnly*” refers to the model with only global max-pooling. Bold numbers are the best results among all presented models. We found that **1)** inputting correct words to an autoregressive decoder is not necessary; **2)** predict-all-words decoders work roughly the same as autoregressive decoders; **3)** mean+max pooling provides stronger transferability than the max-pooling alone does. The table supports our choice of the predict-all-words CNN decoder and the way of producing vector representations from the bi-directional RNN encoder.

Finding I: It is not necessary to input the correct words into an autoregressive decoder for learning sentence representations.

The experimental design was inspired by Bengio et al. (2015). The model we designed for the experiment has a bi-directional GRU as the encoder, and an autoregressive decoder, including both RNN and CNN. We started by analysing the effect of different sampling strategies of the input words on learning an auto-regressive decoder.

We compared three sampling strategies of input words in decoding the target sequence with an autoregressive decoder: (1) **Teacher-Forcing**: the decoder always gets the ground-truth words; (2) **Always Sampling**: at time step t , a word is sampled from the multinomial distribution predicted at time step $t - 1$; (3) **Uniform Sampling**: a word is uniformly sampled from the dictionary V , then fed to the decoder at every time step.

The results are presented in Table 1 (top two subparts). As we can see, the three decoding settings do not differ significantly in terms of the performance on selected downstream tasks, with RNN or CNN as the decoder. The results show that, in terms of learning good sentence representations,

the autoregressive decoder doesn’t require the correct ground-truth words as the inputs.

Finding II: The model with an autoregressive decoder performs similarly to the model with a predict-all-words decoder.

With **Finding I**, we conducted an experiment to test whether the model needs an autoregressive decoder at all. In this experiment, the goal is to compare the performance of the predict-all-words decoders and that of the autoregressive decoders separate from the RNN/CNN distinction, thus we designed a predict-all-words CNN decoder and RNN decoder. The predict-all-words CNN decoder is described in Section 2, which is a stack of three convolutional layers, and all words are predicted once at the output layer of the decoder. The predict-all-words RNN decoder is built based on our CNN decoder. To keep the number of parameters of the two predict-all-words decoder roughly the same, we replaced the last two convolutional layers with a bidirectional GRU.

The results are also presented in Table 1 (3rd and 4th subparts). The performance of the predict-all-words RNN decoder does not significantly differ from that of any one of the autoregressive RNN de-

coders, and the same situation can be also observed in CNN decoders.

These two findings indeed support our choice of using a predict-all-words CNN as the decoder, as it brings the model high training efficiency while maintaining strong transferability.

3.2 Mean+Max Pooling

Since the encoder is a bi-directional RNN in our model, we have multiple ways to select/compute on the generated hidden states to produce a sentence representation. Instead of using the last hidden state as the sentence representation as done in skip-thoughts (Kiros et al., 2015) and SDAE (Hill et al., 2016), we followed the idea proposed in Chen et al. (2016). They built a model for supervised training on the SNLI dataset (Bowman et al., 2015) that concatenates the outputs from a global mean pooling over time and a global max pooling over time to serve as the sentence representation, and showed a performance boost on the SNLI task. Conneau et al. (2017) found that the model with global max pooling function provides stronger transferability than the model with a global mean pooling function does.

In our proposed RNN-CNN model, we empirically show that the mean+max pooling provides stronger transferability than the max pooling alone does, and the results are presented in the last two sections of Table 1. The concatenation of a mean-pooling and a max pooling function is actually a parameter-free composition function, and the computation load is negligible compared to all the heavy matrix multiplications in the model. Also, the non-linearity of the max pooling function augments the mean pooling function for constructing a representation that captures a more complex composition of the syntactic information.

3.3 Tying Word Embeddings and Word Prediction Layer

We choose to share the parameters in the word embedding layer of the RNN encoder and the word prediction layer of the CNN decoder. Tying was shown in both Inan et al. (2016) and Press and Wolf (2017), and it generally helped to learn a better language model. In our model, tying also drastically reduces the number of parameters, which could potentially prevent overfitting.

Furthermore, we initialise the word embeddings with pretrained word vectors, such as word2vec (Mikolov et al., 2013) and GloVe (Pennington et al.,

2014), since it has been shown that these pretrained word vectors can serve as a good initialisation for deep learning models, and more likely lead to better results than a random initialisation.

3.4 Study of the Hyperparameters in Our Model Design

We studied hyperparameters in our model design based on three out of 10 downstream tasks, which are SICK-R, SICK-E (Marelli et al., 2014), and STS14 (Agirre et al., 2014). The first model we created, which is reported in Section 2, is a decent design, and the following variations didn't give us much performance change except improvements brought by increasing the dimensionality of the encoder. However, we think it is worth mentioning the effect of hyperparameters in our model design. We present the Table 1 in the supplementary material and we summarise it as follows:

1. Decoding the next sentence performed similarly to decoding the subsequent contiguous words.
2. Decoding the subsequent 30 words, which was adopted from the skip-thought training code³, gave reasonably good performance. More words for decoding didn't give us a significant performance gain, and took longer to train.
3. Adding more layers into the decoder and enlarging the dimension of the convolutional layers indeed slightly improved the performance on the three downstream tasks, but as training efficiency is one of our main concerns, it wasn't worth sacrificing training efficiency for the minor performance gain.
4. Increasing the dimensionality of the RNN encoder improved the model performance, and the additional training time required was less than needed for increasing the complexity in the CNN decoder. We report results from both smallest and largest models in Table 2.

4 Experiment Settings

The vocabulary for unsupervised training contains the 20k most frequent words in BookCorpus. In order to generalise the model trained with a relatively small, fixed vocabulary to the much larger set of all possible English words, we followed the vocabulary expansion method proposed in Kiros et al. (2015), which learns a linear mapping from the pretrained word vectors to the learnt RNN word

³<https://github.com/ryankiros/skip-thoughts/>

Model	Hrs	SICK-R	SICK-E	STS14	MSRP	TREC MR	CR	SUBJ	MPQA	SST	
Measurement		<i>r</i>	Acc.	<i>r/p</i>	Acc./F1	Accuracy					
<i>Unsupervised training with unordered sentences</i>											
ParagraphVec	4	-	-	0.42/0.43	72.9/81.1	59.4	60.2	66.9	76.3	70.7	-
word2vec BOW	2	0.8030	78.7	0.65/0.64	72.5/81.4	83.6	77.7	79.8	90.9	88.3	79.7
fastText BOW	-	0.8000	77.9	0.63/0.62	72.4/81.2	81.8	76.5	78.9	91.6	87.4	78.8
SIF (GloVe+WR)	-	0.8603	84.6	0.69/-	-/-	-	-	-	-	-	82.2
GloVe BOW	-	0.8000	78.6	0.54/0.56	72.1/80.9	83.6	78.7	78.5	91.6	87.6	79.8
SDAE	72	-	-	0.37/0.38	73.7/80.7	78.4	74.6	78.0	90.8	86.9	-
<i>Unsupervised training with ordered sentences - BookCorpus</i>											
FastSent	2	-	-	0.63/0.64	72.2/80.3	76.8	70.8	78.4	88.7	80.6	-
FastSent+AE	2	-	-	0.62/0.62	71.2/79.1	80.4	71.8	76.5	88.8	81.5	-
Skip-thoughts	336	0.8580	82.3	0.29/0.35	73.0/82.0	92.2	76.5	80.1	93.6	87.1	82.0
Skip-thought+LN	720	0.8580	79.5	0.44/0.45	-	88.4	79.4	83.1	93.7	89.3	82.9
combine CNN-LSTM	-	0.8618	-	-	76.5/83.8	92.6	77.8	82.1	93.6	89.4	-
<i>small RNN-CNN</i> †	20	0.8530	82.6	0.58/0.56	75.6/82.9	89.2	77.6	80.3	92.3	87.8	82.8
<i>large RNN-CNN</i> †	34	0.8698	85.2	0.59/0.57	75.1/83.2	92.2	79.7	81.9	94.0	88.7	84.1
<i>Unsupervised training with ordered sentences - Amazon Book Review</i>											
<i>small RNN-CNN</i> †	21	0.8476	82.7	0.53/0.53	73.8/81.5	84.8	83.3	83.0	94.7	88.2	87.8
<i>large RNN-CNN</i> †	33	0.8616	84.3	0.51/0.51	75.7/82.8	90.8	85.3	86.8	95.3	89.0	88.3
<i>Unsupervised training with ordered sentences - Amazon Review</i>											
BYTE m-LSTM	720	0.7920	-	-	75.0/ 82.8	-	86.9	91.4	94.6	88.5	-
<i>Supervised training - Transfer learning</i>											
DiscSent	8	-	-	-	75.0/-	87.2	-	-	93.0	-	-
DisSent Books 8	-	0.8170	81.5	-	-/-	87.2	82.9	81.4	93.2	90.0	80.2
CaptionRep BOW	24	-	-	0.46/0.42	-	72.2	61.9	69.3	77.4	70.8	-
DictRep BOW	24	-	-	0.67/ 0.70	68.4/76.8	81.0	76.7	78.7	90.7	87.2	-
InferSent(SNLI)	<24	0.8850	84.6	0.68/0.65	75.1/82.3	88.7	79.9	84.6	92.1	89.8	83.3
InferSent(AllNLI)	<24	0.8840	86.3	0.70/0.67	76.2/83.1	88.2	81.1	86.3	92.4	90.2	84.6

Table 2: **Related Work and Comparison.** As presented, our designed asymmetric RNN-CNN model has strong transferability, and is overall better than existing unsupervised models in terms of fast training speed and good performance on evaluation tasks. “†”s refer to our models, and “small/large” refers to the dimension of representation as 1200/4800. **Bold** numbers are the best ones among the models with same training and transferring setting, and underlined numbers are best results among all transfer learning models. The training time of each model was collected from the paper that proposed it.

vectors. Thus, the model benefits from the generalisation ability of the pretrained word embeddings.

The downstream tasks for evaluation include semantic relatedness (SICK, Marelli et al. (2014)), paraphrase detection (MSRP, Dolan et al. (2004)), question-type classification (TREC, Li and Roth (2002)), and five benchmark sentiment and subjective datasets, which include movie review sentiment (MR, Pang and Lee (2005)), SST, Socher et al. (2013)), customer product reviews (CR, Hu and Liu (2004)), subjectivity/objectivity classification (SUBJ, Pang and Lee (2004)), opinion polarity (MPQA, Wiebe et al. (2005)), semantic textual similarity (STS14, Agirre et al. (2014)), and SNLI (Bowman et al., 2015). After unsupervised training, the encoder is fixed, and applied as a representation

extractor on the 10 tasks.

To compare the effect of different corpora, we also trained two models on Amazon Book Review dataset (without ratings) which is the largest subset of the Amazon Review dataset (McAuley et al., 2015) with 142 million sentences, about twice as large as BookCorpus.

Both training and evaluation of our models were conducted in PyTorch⁴, and we used SentEval⁵ provided by Conneau et al. (2017) to evaluate the transferability of our models. All the models were trained for the same number of iterations with the same batch size, and the performance was measured at the end of training for each of the models.

⁴<http://pytorch.org/>

⁵<https://github.com/facebookresearch/SentEval>

5 Related work and Comparison

Table 2 presents the results on 9 evaluation tasks of our proposed RNN-CNN models, and related work. The “**small RNN-CNN**” refers to the model with the dimension of representation as 1200, and the “**large RNN-CNN**” refers to that as 4800. The results of our “large RNN-CNN” model on SNLI is presented in Table 3.

Model	SNLI (Acc %)
<i>Unsupervised Transfer Learning</i>	
Skip-thoughts (Vendrov et al.)	81.5
large RNN-CNN <i>BookCorpus</i>	81.7
large RNN-CNN <i>Amazon</i>	81.5
<i>Supervised Training</i>	
ESIM (Chen et al.)	86.7
DIIN (Gong et al.)	88.9

Table 3: We implemented the same classifier as mentioned in Vendrov et al. (2015) on top of the features computed by our model. Our proposed RNN-CNN model gets similar result on SNLI as skip-thoughts, but with much less training time.

Our work was inspired by analysing the **skip-thoughts** model (Kiros et al., 2015). The skip-thoughts model successfully applied this form of learning from the context information into unsupervised representation learning for sentences, and then, Ba et al. (2016) augmented the LSTM with proposed layer-normalisation (**Skip-thought+LN**), which improved the skip-thoughts model generally on downstream tasks. In contrast, Hill et al. (2016) proposed the **FastSent** model which only learns source and target word embeddings and is an adaptation of **Skip-gram** (Mikolov et al., 2013) to sentence-level learning without word order information. Gan et al. (2017) applied a CNN as the encoder, but still applied LSTMs for decoding the adjacent sentences, which is called **CNN-LSTM**.

Our **RNN-CNN** model falls in the same category as it is an encoder-decoder model. Instead of decoding the surrounding two sentences as in skip-thoughts, FastSent and the compositional CNN-LSTM, our model only decodes the subsequent sequence with a fixed length. Compared with the hierarchical CNN-LSTM, our model showed that, with a proper model design, the context information from the subsequent words is sufficient for learning sentence representations. Particularly, our proposed **small RNN-CNN** model runs roughly three

times faster than our implemented skip-thoughts model⁶ on the same GPU machine during training.

Proposed by Radford et al. (2017), **BYTE m-LSTM** model uses a multiplicative LSTM unit (Krause et al., 2016) to learn a language model. This model is simple, providing next-byte prediction, but achieves good results likely due to the extremely large training corpus (Amazon Review data, McAuley et al. (2015)) that is also highly related to many of the sentiment analysis downstream tasks (domain matching).

We experimented with the Amazon Book review dataset, the largest subset of the Amazon Review. This subset is significantly smaller than the full Amazon Review dataset but twice as large as BookCorpus. Our RNN-CNN model trained on the Amazon Book review dataset resulted in performance improvement on all single-sentence classification tasks relative to that achieved with training under BookCorpus.

Unordered sentences are also useful for learning representations of sentences. **ParagraphVec** (Le and Mikolov, 2014) learns a fixed-dimension vector for each sentence by predicting the words within the given sentence. However, after training, the representation for a new sentence is hard to derive, since it requires optimising the sentence representation towards an objective. **SDAE** (Hill et al., 2016) learns the sentence representations with a denoising auto-encoder model. Our proposed RNN-CNN model trains faster than SDAE does, and also because we utilised the sentence-level continuity as a supervision which SDAE doesn’t, our model largely performs better than SDAE.

Another transfer approach is to learn a supervised discriminative classifier by distinguishing whether the sentence pair or triple comes from the same context. Li and Hovy (2014) proposed a model that learns to classify whether the input sentence triplet contains three contiguous sentences. **DiscSent** (Jernite et al., 2017) and **DisSent** (Nie et al., 2017) both utilise the annotated explicit discourse relations, which is also good for learning sentence representations. It is a very promising research direction since the proposed models are generally computational efficient and have clear intuition, yet more investigations need to be done to augment the performance.

Supervised training for transfer learning is

⁶We reimplemented the skip-thoughts model in PyTorch, and it took roughly four days to train for only one epoch on BookCorpus.

also promising when a large amount of human-annotated data is accessible. [Conneau et al. \(2017\)](#) proposed the **InferSent** model, which applies a bi-directional LSTM as the sentence encoder with multiple fully-connected layers to classify whether the hypothesis sentence entails the premise sentence in SNLI ([Bowman et al., 2015](#)), and MultiNLI ([Williams et al., 2017](#)). The trained model demonstrates a very impressive transferability on downstream tasks, including both supervised and unsupervised. Our RNN-CNN model trained on Amazon Book Review data in an unsupervised way has better results on supervised tasks than **InferSent** but slightly inferior results on semantic relatedness tasks. We argue that labelling a large amount of training data is time-consuming and costly, while unsupervised learning provides great performance at a fraction of the cost. It could potentially be leveraged to initialise or more generally augment the costly human labelling, and make the overall system less costly and more efficient.

6 Discussion

In [Hill et al. \(2016\)](#), internal consistency is measured on five single sentence classification tasks (MR, CR, SUBJ, MPQA, TREC), MSRP and STS-14, and was found to be only above the “acceptable” threshold. They empirically showed that models that worked well on supervised evaluation tasks generally didn’t perform well on unsupervised ones. This implies that we should consider supervised and unsupervised evaluations separately, since each group has higher internal consistency.

As presented in [Table 2](#), the encoders that only sum over pretrained word vectors perform better overall than those with RNNs on unsupervised evaluation tasks, including STS14. In recent proposed log-bilinear models, such as **FastSent** ([Hill et al., 2016](#)) and **SiameseBOW** ([Kenter et al., 2016](#)), the sentence representation is composed by summing over all word representations, and the only tunable parameters in the models are word vectors. These resulting models perform very well on unsupervised tasks. By augmenting the pretrained word vectors with a weighted averaging process, and removing the top few principal components, which mainly encode frequently-used words, as proposed in [Arora et al. \(2017\)](#) and [Mu et al. \(2018\)](#), the performance on the unsupervised evaluation tasks gets even better. Prior work suggests that incorporating word-level information helps the model to perform

better on cosine distance based semantic textual similarity tasks.

Our model predicts all words in the target sequence at once, without an autoregressive process, and ties the word embedding layer in the encoder with the prediction layer in the decoder, which explicitly uses the word vectors in the target sequence as the supervision in training. Thus, our model incorporates the word-level information by using word vectors as the targets, and it improves the model performance on STS14 compared to other RNN-based encoders.

[Arora et al. \(2017\)](#) conducted an experiment to show that the word order information is crucial in getting better results on supervised tasks. In our model, the encoder is still an RNN, which explicitly utilises the word order information. We believe that the combination of encoding a sentence with its word order information and decoding all words in a sentence independently inherently leverages the benefits from both log-linear models and RNN-based models.

7 Conclusion

Inspired by learning to exploit the contextual information present in adjacent sentences, we proposed an asymmetric encoder-decoder model with a suite of techniques for improving context-based unsupervised sentence representation learning. Since we believe that a simple model will be faster in training and easier to analyse, we opt to use simple techniques in our proposed model, including 1) an RNN as the encoder, and a predict-all-words CNN as the decoder, 2) learning by inferring subsequent contiguous words, 3) mean+max pooling, and 4) tying word vectors with word prediction. With thorough discussion and extensive evaluation, we justify our decision making for each component in our RNN-CNN model. In terms of the performance and the efficiency of training, we justify that our model is a fast and simple algorithm for learning generic sentence representations from unlabelled corpora. Further research will focus on how to maximise the utility of the context information, and how to design simple architectures to best make use of it.

References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel M. Cer, Mona T. Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. Semeval-2014 task 10: Multilingual semantic textual similarity. In *SemEval@COLING*.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *ICLR*.
- Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *CoRR*, abs/1607.06450.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *EMNLP*.
- Qian Chen, Xiaodan Zhu, Zhenhua Ling, Si Wei, and Hui Jiang. 2016. Enhancing and combining sequential and tree lstm for natural language inference. *arXiv preprint arXiv:1609.06038*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *EMNLP*.
- William B. Dolan, Chris Quirk, and Chris Brockett. 2004. Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources. In *COLING*.
- J. R. Firth. 1957. A synopsis of linguistic theory.
- Zhe Gan, Yunchen Pu, Ricardo Henao, Chunyuan Li, Xiaodong He, and Lawrence Carin. 2017. Learning generic sentence representations using convolutional neural networks. In *EMNLP*.
- Yichen Gong, Heng Luo, and Jian Zhang. 2017. Natural language inference over interaction space. *CoRR*, abs/1709.04348.
- Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. Learning distributed representations of sentences from unlabelled data. In *HLT-NAACL*.
- Sepp Hochreiter and Juergen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9:1735–1780.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *KDD*.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2016. Tying word vectors and word classifiers: A loss framework for language modeling. *CoRR*, abs/1611.01462.
- Yacine Jernite, Samuel R. Bowman, and David Sonntag. 2017. Discourse-based objectives for fast unsupervised sentence representation learning. *CoRR*, abs/1705.00557.
- Tom Kenter, Alexey Borisov, and Maarten de Rijke. 2016. Siamese cbow: Optimizing word embeddings for sentence representations. In *ACL*.
- Jamie Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *NIPS*.
- Ben Krause, Liang Lu, Iain Murray, and Steve Renals. 2016. Multiplicative lstm for sequence modelling. *CoRR*, abs/1609.07959.
- Quoc V. Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *ICML*.
- Jiwei Li and Eduard H. Hovy. 2014. A model of coherence based on distributed sentence representation. In *EMNLP*.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *COLING*.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*.
- Julian J. McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-based recommendations on styles and substitutes. In *SIGIR*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- Jiaqi Mu, Suma Bhat, and Pramod Viswanath. 2018. All-but-the-top: Simple and effective postprocessing for word representations. In *ICLR*.
- Allen Nie, Erin D. Bennett, and Noah D. Goodman. 2017. Dissent: Sentence representation learning from explicit discourse relations. *CoRR*, abs/1710.04334.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL*.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *ACL*.

- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Ofir Press and Lior Wolf. 2017. Using the output embedding to improve language models. In *EACL*.
- Alec Radford, Rafal Józefowicz, and Ilya Sutskever. 2017. Learning to generate reviews and discovering sentiment. *CoRR*, abs/1704.01444.
- 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 *EMNLP*.
- Shuai Tang, Hailin Jin, Chen Fang, Zhaowen Wang, and Virginia R. de Sa. 2017. Rethinking skip-thought: A neighborhood based approach. In *RepLANLP, ACL Workshop*.
- Peter D. Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *J. Artif. Intell. Res.*, 37:141–188.
- Ivan Vendrov, Jamie Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2015. Order-embeddings of images and language. *CoRR*, abs/1511.06361.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39:165–210.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *CoRR*, abs/1704.05426.