

---

# Paraphrase Thought: Sentence Embedding Module Imitating Human Language Recognition

---

Myeongjun Jang<sup>1</sup> Pilsung Kang<sup>1</sup>

## Abstract

Sentence embedding is an important research topic in natural language processing. It is essential to generate a good embedding vector that fully reflects the semantic meaning of a sentence in order to achieve an enhanced performance for various natural language processing tasks, such as machine translation and document classification. Thus far, various sentence embedding models have been proposed, and their feasibility has been demonstrated through good performances on tasks following embedding, such as sentiment analysis and sentence classification. However, because the performances of sentence classification and sentiment analysis can be enhanced by using a simple sentence representation method, it is not sufficient to claim that these models fully reflect the meanings of sentences based on good performances for such tasks. In this paper, inspired by human language recognition, we propose the following concept of semantic coherence, which should be satisfied for a good sentence embedding method: similar sentences should be located close to each other in the embedding space. Then, we propose the Paraphrase-Thought (P-thought) model to pursue semantic coherence as much as possible. Experimental results on two paraphrase identification datasets (MS COCO and STS benchmark) show that the P-thought models outperform the benchmarked sentence embedding methods.

**Keywords:** *Sentence embedding, Recurrent neural network, Paraphrase, Semantic coherence, Natural language processing*

## 1. Introduction

Sentence embedding, which transforms sentences into low-dimensional vector values reflecting their meanings, is a highly important task in natural language processing (NLP). By mapping unstructured text data into a certain form of structured representation, the embedding vector can enhance the performances of various NLP tasks, such as machine translation (Artetxe et al., 2017; Lee et al., 2016; Zhao & Zhang, 2016), document classification (Conneau et al., 2017b; Zhou et al., 2016), and sentence matching (Wan et al., 2016). As sentence embedding plays an import role in NLP, various methods (Kiros et al., 2015; Pagliardini et al., 2017; Hill et al., 2016; Arora et al., 2017; Conneau et al., 2017a; Chen, 2017) have been proposed since the advent of the Doc2vec method (Le & Mikolov, 2014). Typically, these methods exhibit better performances than benchmarked embedding methods for common NLP tasks, such as document classification or sentiment analysis. However, this is not a direct evaluation of how well semantic meanings are preserved by the proposed embedding method.

Indirect methods for evaluating sentence embedding are not sufficient to evaluate the main property of sentence embedding techniques, i.e., how well semantic relationships between sentences are preserved. Iyyer et al. (2015) showed that it is possible to achieve a fairly good performance in document classification using a simple document representation vector, i.e., an average of word vectors in the document. Even for classic document representation methods, in which word sequences or semantic relationships between words are not considered, e.g., bag of words (BoW) or term frequency-inverse document frequency (TF-IDF), highly accurate classification results can be achieved using a Naïve Bayesian classifier (Soumya George & Joseph, 2014). This means that a good performance on a classification task can be achieved without the use of embedding vectors. In other words, a good classification performance for common NLP tasks using a certain type of sentence embedding method does not guarantee that the embedding method can successfully

---

<sup>1</sup>School of Industrial Management Engineering, Korea University, Seoul, South Korea. Correspondence to: Myeongjun Jang <xkxpa@korea.ac.kr>, Pilsung Kang <pilsung\_kang@korea.ac.kr>.

preserve the semantic relationship between sentences.

In this paper, in order to overcome the limitations of indirect sentence embedding evaluation strategies, we propose the following concept of semantic coherence, which should be satisfied by a good sentence embedding method: sentences having similar meanings should be placed close to each other in the embedding space. Then, we propose a new sentence embedding model named Paraphrase-Thought (P-thought), which can maximally pursue semantic coherence during training. The P-thought model is designed as a dual generation model, which receives a single sentence as input and generates both the input sentence and its paraphrase sentence simultaneously. The proposed P-thought model is evaluated through a task of measuring the semantic coherence and the STS Benchmark task. Experimental results show that the proposed P-thought model yields a better performance than benchmarked models in both tasks.

The remainder of this paper is organized as follows. In Section 2, we briefly review previous research on sentence embedding. In Section 3, we propose the concept of semantic coherence and a new metric: paraphrase coherence (P-coherence). In Section 4, we describe the structure of the P-thought model. In Section 5, experimental settings are described for each task, followed by results and discussions. In Section 6, we conclude the present work with some discussion of future research directions.

## 2. Related work

Recent work on sentence embedding ranges from simple extensions of the word embedding vector (Le & Mikolov, 2014; Arora et al., 2017; Pagliardini et al., 2017; Chen, 2017; Wieting et al., 2015) to neural network models specialized for handling a sequence of words appearing in a sentence (Kiros et al., 2015; Conneau et al., 2017a). The Distributed Bag of Words version of Paragraph Vector (PV-DBOW) and Distributed Memory Model of Paragraph Vectors (PV-DM) methods, which were proposed in Doc2vec (Le & Mikolov, 2014), learn sentence vectors based on the same principle: maximizing the probability to predict words in the same sentence. Arora et al. (2017) proposed a model that computes the sentence embedding vector as a weighted average of word embedding vectors in a sentence. By re-weighting the weights of words in a sentence, the authors achieved an improved performance in textual similarity tasks, and outperformed a complex model based on recurrent neural network (RNN). Unlike in Doc2vec (Le & Mikolov, 2014), in Doc2vecC (Chen, 2017), the sentence embed-

ding vector is defined as a simple average of word embedding vectors. The idea behind doc2vecC, i.e., using an average of word embedding vectors to represent the global context of the sentence, had already been proposed by Huang et al. (2012). In addition, doc2vecC applies a corruption mechanism that randomly removes words from a sentence and generates a sentence embedding vector with the remaining words. This simple idea significantly reduced the total amount of training time. Similar to previous methods, Sent2vec (Pagliardini et al., 2017) defines the sentence vector as an average of word embedding vectors. However, unlike other models using word embedding vectors of single words (i.e., uni-gram), it considers n-gram vectors in addition to uni-gram vectors when training the sentence embedding model.

The Skip-thought model (Kiros et al., 2015), which has a sequence to sequence (Seq2Seq) structure, is an extension of the Skip-gram (Mikolov et al., 2013b) model, where the basic unit for network learning is a sentence instead of a word. Similar to the Skip-gram model, which learns word embedding vectors by training the network to predict the surrounding words when the center word is given, Skip-thought is trained to encode the input sentence and generate its preceding and following sentences. By using the generated sentence vectors as the input of a simple linear model, Skip-thought exhibited an improved performance for document classification and sentiment analysis. Inspired by previous results in computer vision, where many models are pretrained based on ImageNet (Deng et al., 2009), Conneau et al. (2017a) conducted research on whether supervised learning tasks are helpful for learning sentence embedding vectors. Through experiments, Conneau et al. (2017a) claimed that sentence embedding vectors generated from a model that is trained based on a natural language inference (NLI) task yield a state-of-the-art performance when leveraged in other NLP tasks. In particular, they found that a model with a bi-directional Long-short term memory (LSTM) structure and max pooling trained on the Stanford Natural Language Inference (SNLI) dataset (Bowman et al., 2015), named InferSent, exhibited the best performance.

## 3. Semantic coherence

### 3.1. Defining semantic coherence

Although two sentences may employ different words or different structures, people will recognize them as the same sentence as long as the implied semantic meanings are highly similar. Consider the following two sentences:

- **Sentence 1:** Jang was caught by professor Kang while playing the computer game in the lab.
- **Sentence 2:** Professor Kang came to the lab and witnessed Jang playing the computer game.

Although these two sentences exhibit a clear difference with respect to both the sentence structure and word usage, people can immediately perceive that they convey the same meaning. Hence, a good sentence embedding approach should satisfy the property that if two sentences have different structures but convey the same meaning (i.e., paraphrase sentences), then they should have the same, or at least similar, embedding vectors. Based on this, we define semantic coherence as follows.

**Definition 1.** The degree of semantic coherence of a sentence embedding model is proportional to the similarity between the representation vectors of paraphrase sentences generated by the model.

If the representation vectors of paraphrase sentences are located close to each other in the embedding space, this implies that there is little difference between their vector values. Thus, when the representation vector value of a sentence is given, it should be possible to generate the given sentence and its paraphrase sentences. Consequently, we can derive the following hypothesis.

**Hypothesis 1.** If it is possible to generate an input sentence and its paraphrase sentence simultaneously from the vector value of the input sentence, then the sentence embedding model can enhance the semantic coherence.

In this study, we propose a new sentence embedding model to satisfy the above hypothesis.

### 3.2. Evaluating semantic coherence: paraphrase coherence

To evaluate semantic coherence, we should measure the densities of paraphrase sentences. This requires multiple pairs of paraphrase sentences that share the same meaning. Thus, previous metrics that simply calculate the matching degree of two sentences are insufficient.

In this study, inspired by topic coherence, which is used to determine the optimal number of topics in topic modeling, we propose a new evaluation metric

called paraphrase coherence (P-coherence) to measure the semantic coherence. Topic coherence measures how effectively the highly weighted top  $k$  words of a topic satisfy coherence (Newman et al., 2010); topic coherence is computed as follows:

$$Topic - coherence = \sum_{i < j} Score(w_i, w_j), \quad (1)$$

where  $w_i$  and  $w_j$  are the top  $i^{th}$  and  $j^{th}$  words of the same topic, respectively. Although various methods exist to define the score between two words (Röder et al., 2015), we adopted the idea of the pointwise mutual information (PMI) measure (Newman et al., 2010), defined as follows:

$$Score_{PMI}(w_i, w_j) = \log \frac{p(w_i, w_j)}{p(w_i)p(w_j)}, \quad (2)$$

where  $p(w_i, w_j)$  is the probability of words  $w_i$  and  $w_j$  appearing together in a randomly selected document, and  $p(w_i)$  and  $p(w_j)$  are the marginal probabilities that the words  $w_i$  and  $w_j$  appear in the randomly selected document, respectively.

Unlike in topic coherence, which defines the probability of two words appearing together based on a simple word count, we should consider the relationship between two sentences by leveraging their representation vector values. Hence, we replace the co-occurrence probability  $p(w_i, w_j)$  in topic coherence with the dot product of two sentence representation vectors because the dot product of two vectors is widely used as an unnormalized probability in many studies (Karpathy et al., 2014; Karpathy & Fei-Fei, 2015). Next, we replace the marginal probability for word occurrence in topic coherence with the  $L_2$ -norm of the sentence vector, derived from the dot product. As a result, the score between two sentences takes a value between 0 and 1: the higher the score value, the stronger is the relationship between the two sentences. The equation representing the proposed score is as follows:

$$Score(sv_i, sv_j) = \frac{sv_i \cdot sv_j}{\|sv_i\| \|sv_j\|}, \quad (3)$$

where  $sv_i$  and  $sv_j$  are the representation vectors of the sentences  $i$  and  $j$ , respectively. Finally, P-coherence is defined as the average score of all pairs of paraphrase sentences:

$$P - coherence(U_k) = Average\left(\frac{sv_i \cdot sv_j}{\|sv_i\| \|sv_j\|}\right), \quad (4)$$

$$sv_i, sv_j \in U_k, k = 1, \dots, N,$$

where  $U_k$  is the  $k^{th}$  set of paraphrase sentences, and  $N$  is the number of paraphrase sets. For instance, if there

are four paraphrase sentences for each paraphrase set, then the P-coherence for each paraphrase set is calculated as the average score of  ${}_4C_2 = 6$  sentence pairs. The total P-coherence is the average P-coherence for each paraphrase set:

$$P - coherence_{Total} = \frac{1}{N} \sum_{k=1}^N P - coherence(U_k). \quad (5)$$

## 4. Paraphrase thought

### 4.1. Model structure

Assume that a sentence tuple  $(s, p)$  is given, where  $p$  is the paraphrase sentence of the sentence  $s$ . Let  $x_t$  be the  $t^{th}$  word of the sentence  $s$  and  $y_t$  be the  $t^{th}$  word of the sentence  $p$ . To maximize the semantic coherence defined above, it should be possible to generate both the sentence itself and its paraphrase sentence from the representation vector of an input sentence. Therefore, the proposed P-thought model is designed as a dual generation model, which generates both  $s$  and  $p$  simultaneously when the sentence tuple  $(s, p)$  is given.

We employed an Seq2Seq structure with a gated recurrent unit (GRU) (Cho et al., 2014) cell for the P-thought model. The *encoder* transforms the sequence of words of an input sentence into a fixed-sized representation vector, whereas the *decoder* generates the target sentence based on the given sentence representation vector. The proposed P-thought model has two decoders. When the input sentence is given, the first decoder, named *auto-decoder*, generates the input sentence as it is. The second decoder, named *paraphrase-decoder*, generates the paraphrase sentence of the input sentence.

### 4.2. Objective

Similar to other sequence learning tasks in NLP, the purpose of the P-thought model is to minimize the negative log likelihoods of the two decoders. Furthermore, according to Hypothesis 1, the P-thought model should satisfy the condition that it can encode the sentence  $s$  and generate the sentences  $s$  and  $p$  simultaneously when the sentence pair  $(s, p)$  is given. This condition can be written as follows:

$$P(s)P(s|s; \theta_{ss}) = P(s)P(p|s; \theta_{sp}), \quad (6)$$

where  $P(s)$  is the marginal probability of the input sentence  $s$ , and  $\theta_{ss}$  and  $\theta_{sp}$  are the parameters of *auto-decoder* and *paraphrase-decoder*, respectively. Thus, similarly to the work of Xia et al. (2017), the

problem can be formulated as the following multi-objective optimization problem:

$$\begin{aligned} \textbf{Objective 1} : & \min_{\theta_{ss}} l_A(f(s; \theta_{ss}), s), \\ \textbf{Objective 2} : & \min_{\theta_{sp}} l_P(f(s; \theta_{sp}), p), \end{aligned} \quad (7)$$

$$s.t \quad P(s)P(s|s; \theta_{ss}) = P(s)P(p|s; \theta_{sp}),$$

where  $l_A$  and  $l_P$  are the negative log likelihoods of *auto-decoder* and *paraphrase-decoder*, respectively. In this case, the constraint term can be rewritten as follows:

$$-\log P(s|s; \theta_{ss}) = -\log P(p|s; \theta_{sp}), \quad (8)$$

The left and right terms of transformed equation represent the negative log likelihoods of *auto-decoder* and *paraphrase-decoder*, respectively. Hence, the constraint term can be written as follows:

$$l_A(f(s; \theta_{ss}), s) = l_P(f(s; \theta_{sp}), p). \quad (9)$$

By introducing the Lagrange multiplier, the multi-objective optimization problem is transformed into the following minimization problem:

$$\begin{aligned} \min L = & l_A(f(s; \theta_{ss}), s) + l_P(f(s; \theta_{sp}), p) \\ & - \lambda(l_A(f(s; \theta_{ss}), s) - l_P(f(s; \theta_{sp}), p)) \\ = & (1 - \lambda)l_A(f(s; \theta_{ss}), s) + (1 + \lambda)l_P(f(s; \theta_{sp}), p), \\ \text{where } & \lambda \neq 0. \end{aligned} \quad (10)$$

In this case, a value of  $\lambda > 1$  or  $\lambda < -1$  leads to maximizing the negative log likelihood of *auto-decoder* and that of *paraphrase-decoder*, respectively. To avoid this problem, the allowable range for  $\lambda$  is set to  $-1 < \lambda < 0$  or  $0 < \lambda < 1$ . Under this condition, minimizing  $L$  is equivalent to minimizing  $L'$ :

$$\begin{aligned} \min L' = & l_A(f(s; \theta_{ss}), s) + \alpha l_P(f(s; \theta_{sp}), p), \\ \text{where } & \alpha = \frac{1 + \lambda}{1 - \lambda}, \end{aligned} \quad (11)$$

where  $\alpha$  is the hyperparameter of the P-thought model. This should be greater than 1 or be the value between 0 and 1, because  $-1 < \lambda < 0$  and  $0 < \lambda < 1$ . However, it is desirable to set the appropriate  $\alpha$ -value to greater than 1, considering that auto decoding is trivial copying task which is much easier than paraphrase generation. Experimental results also demonstrated that the performance is degraded for small  $\alpha$  value. Thus, the objective of the P-thought model is the sum of the negative log likelihood of *auto-decoder* with that of *paraphrase-decoder* with a higher weight:

$$\begin{aligned} \text{Loss} = & l_{auto}(f(s; \theta_{ss}), s) + \alpha l_{para}(f(s; \theta_{sp}), p), \\ \text{where } & \alpha > 1. \end{aligned} \quad (12)$$

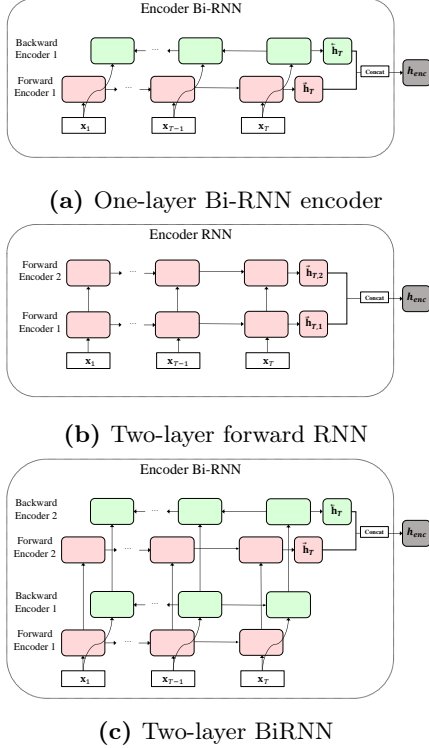


Figure 1. Three encoder structures of the P-thought model

### 4.3. Vocabulary expansion

The number of unique words appearing in our training dataset is only about 35,000, which is considerably fewer than the number of words in the English language. This may be problematic, in that many words are treated as out of vocabulary after model training. To solve this problem, motivated by the idea of cross-lingual embedding (Mikolov et al., 2013a), Skip-thought attempts to learn a matrix that maps the words of a pretrained word2vec model (Mikolov et al., 2013b) to one of 20,000 words in their training dataset. However, this approach suffers from the problem that a word can be mapped to another word whose actual meaning is significantly different, only because it has a high similarity with the original word in the embedding space. For example, the word 'endogenous' was mapped to the word 'neuronal,' despite the semantic differences.

We extracted the vector values of words that appear in our training dataset from the pretrained Glove vector (Pennington et al., 2014) to resolve the problem described above. In the pretrained Glove vectors, the semantic relationships between words are reflected in the geometrical structures between word vectors. Therefore, even when vector values of words that are un-

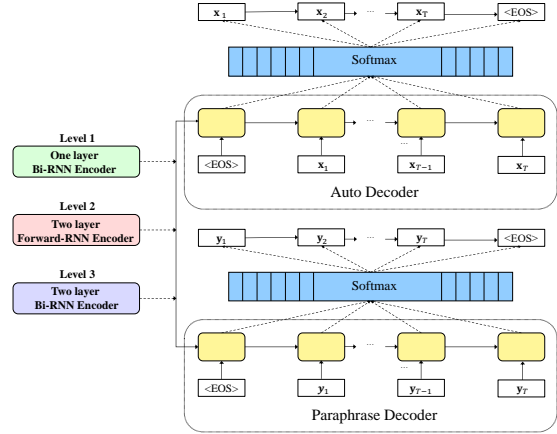


Figure 2. P-thought model structure

Table 1. Training data description

	2014-Validation	2017-Training	Total
No. of unique images	40,504	118,284	123,287
No. of unique captions	202,654	591,753	593,968
No. of unique sentence pairs	811,426	2,368,926	2,467,293
No. of unique words	-	-	34,826

used during training occur, the information loss can be reduced because the geometric relationships between word vectors are well preserved if the model is sufficiently trained. By using this method, we are able to handle 2.1 million words without the effort of training an extra mapping matrix.

## 5. Experiments

### 5.1. Experimental settings

We used the captions of the MS-COCO dataset (Lin et al., 2014) to train the P-thought model. This dataset has been employed in various paraphrase generation studies (Prakash et al., 2016; Gupta et al., 2017). The MS-COCO dataset has more than five captions for each image, which allows us to generate more than  ${}_5P_2 = 20$  unique sentence pairs. For training, we used the 2014-Validation and 2017-Training datasets. Descriptions of these datasets are provided in Table 1. Simple tokenizing was performed as text preprocessing for the captions.

We employed three different encoder structures, as shown in Figure 1, to investigate the model perfor-

Table 2. 2017-Validation dataset description

No. of unique images	No. of unique captions	No. of unique sentence pair	No. of unique words
5,000	25,014	100,142	8,641

Table 3. Experimental results for evaluating the P-coherence

Model	P-coherence
PV-DBOW (Le & Mikolov, 2014)	0.0099
Uni-skip (Kiros et al., 2015)	0.5328
Bi-skip (Kiros et al., 2015)	0.5155
Combine-skip (Kiros et al., 2015)	0.5209
SIF (Arora et al., 2017)	0.4205
Sent2vec Wiki-uni (Pagliardini et al., 2017)	0.4279
Sent2vec Wiki-bi (Pagliardini et al., 2017)	0.4553
InferSent (Conneau et al., 2017a)	0.7454
P-thought (one layer-Bi RNN)	0.7432
P-thought (two layers-Forward RNN)	<b>0.7899</b>
P-thought (two layers-Bi RNN)	<b>0.9725</b>

mances according to different levels of model complexity. The first encoder structure has one layer with a bi-directional RNN (Bi-RNN). The sentence embedding vector of this encoder structure consists of the concatenated values of the final state values of the forward and backward RNN. The second encoder structure contains two layers, with only a forward RNN. The sentence embedding vector is generated by concatenating the final states of both layers. The third encoder structure contains two layers of Bi-RNN. The sentence embedding vector is generated from the concatenated values of the final states of the second layer’s forward and backward RNNs. The overall structure of the P-thought model, including the decoder part, is illustrated in Figure 2. These three models were trained under the same conditions. The number of hidden units is set to 1,200, which results in 2,400-dimensional sentence embedding vectors after concatenation. We employed Xavier initialization (Glorot & Bengio, 2010), and gradient computations and weight updates were performed with a mini-batch size of 128. All models were trained for four epochs using the Adam optimizer (Kingma & Ba, 2014).

Table 4. Extracted sentences for visualization

#### Group 1 (▲)

- 1) Bedroom scene with a bookcase, blue comforter and window.
- 2) A bedroom with a bookshelf full of books.
- 3) This room has a bed with blue sheets and a large bookcase.
- 4) A bed and a mirror in a small room.
- 5) A bed room with a neatly made bed a window and a book shelf

#### Group 2 (◆)

- 1) A male tennis player in white shorts is playing tennis.
- 2) This woman has just returned a volley in tennis.
- 3) A man holding a tennis racket playing tennis.
- 4) The man balances on one leg after serving a tennis ball.
- 5) Someone playing in a tennis tournament with a crowd looking on.

#### Group 3 (■)

- 1) A woman holding a Hello Kitty phone on her hand.
- 2) A woman holds up her phone in front of her face.
- 3) A woman in white shirt holding up a cellphone.
- 4) A woman checking her cell phone with a hello kitty case.
- 5) The Asian girl is holding her Miss Kitty phone.

#### Group 4 (×)

- 1) A plate of food which includes onions, tomato, lettuce, sauce, fries, and a sandwich.
- 2) A sandwich, french fries, bowl of ketchup, onion slice, lettuce slice, tomato slice, and knife sit on the white plate.
- 3) Partially eaten hamburger on a plate with fries and condiments.
- 4) A grilled chicken sandwich sits beside french fries made with real potatoes.
- 5) A sandwich on a sesame seed bun next to a pile of french fries and a cup of ketchup

#### Group 5 (▼)

- 1) Decorated coffee cup and knife sitting on a patterned surface.
- 2) A large knife is sitting in front of a mug has a skull and crossbones.
- 3) A white mug showing pirate skull and bones and a large knife on a counter top.
- 4) There is a white coffee cup with a skull and bones on it next to a knife.
- 5) A close up of a knife and a cup on a surface

## 5.2. P-coherence

To measure the P-coherence, we used the 2017-Validation dataset from the MS-COCO caption dataset, which has no overlap with the training dataset. A description of the dataset used for evaluating the P-coherence is provided in Table 2. We selected PV-DBOW, Skip-thought, SIF, Sent2vec, and InferSent as benchmark models. In the case of PV-DBOW, we employed the datasets used for both training P-thought and evaluating the P-coherence to learn the sentence vectors. For the remaining models, we used the publicly available pretrained models.

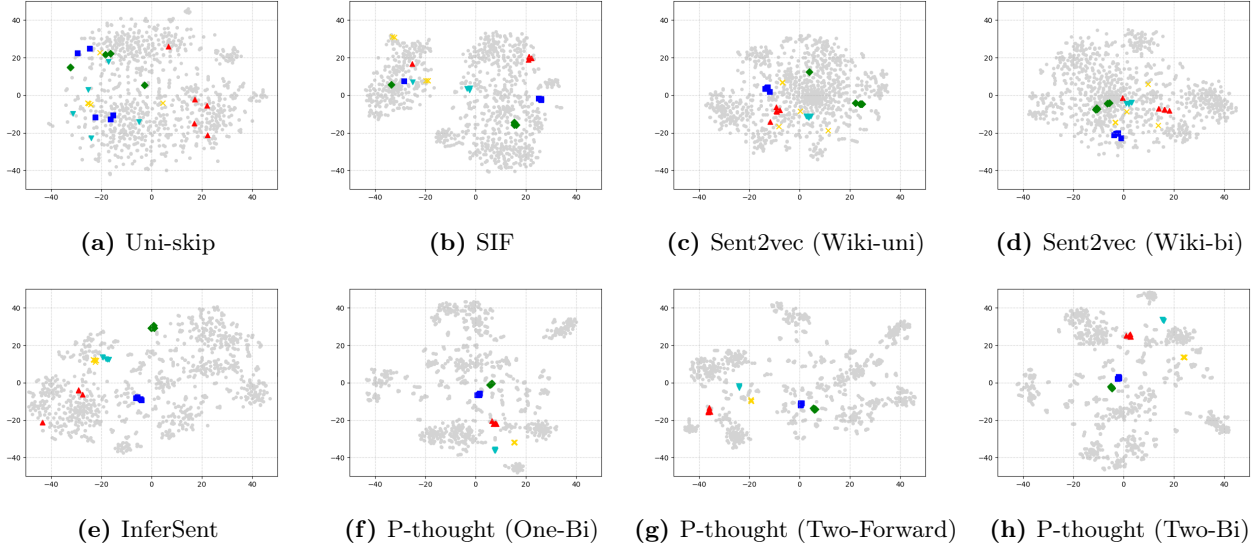


Figure 3. Scatter plots of the five paraphrase sentence groups represented by each sentence embedding method

The experimental results are summarized in Table 3. It can be observed that the P-thought models with relatively complex encoder structures outperformed other benchmarked models. In the case of P-thought with a one-layer Bi-RNN, the P-coherence value is comparable to that of InferSent, and superior to the other benchmarked models. Among the benchmarked models, InferSent yielded a significantly higher P-coherence value than the other models, which implies that InferSent preserved the semantic coherence when learning the sentence representation vectors.

In addition to the quantitative evaluation provided in Table 3, we reduced the generated sentence vectors to two-dimensional vectors using *t*-SNE (Maaten & Hinton, 2008) and created scatter plots to qualitatively investigate how effectively the paraphrase sentences satisfied coherence. For the sake of visualization, we extracted the paraphrase sentences for five images and marked them with different colors and shapes. The extracted paraphrase sentences are presented in Table 4, and the scatter plots are given in Figure 3. It can easily be observed that paraphrase sentence vectors learned by the models with high P-coherence values (P-thought and InferSent) are more concentrated than those of the other models.

### 5.3. STS Benchmark task

We also carried out the STS Benchmark task (Cer et al., 2017) to evaluate how well the models preserve the meanings of sentences through a more generally conducted task. The dataset for this task consists of

Table 5. STS Benchmark task dataset description

-	Train	Dev	Test	Total
# of data	5,749	1,500	1,379	8,628

8,628 sentence pairs and corresponding human rated similarity scores valued between 0 and 5. The purpose of this task is to approximate the similarity scores between sentences based on the embedded vectors. A description of the dataset is summarized in Table 5.

We conducted the experiment in the same manner as for InferSent. For two sentence vectors  $u$  and  $v$ , the component-wise product  $u \cdot v$  and the absolute difference  $|u - v|$  are computed and concatenated to be used as an input. As the target, the human rated similarity score  $y$  is transformed as follows. Let  $r^T = [1, \dots, 5]$  denote a vector that takes integer values between 1 and 5. The target  $y$  is transformed to the distribution  $d$  using the equation below:

$$d_i = \begin{cases} y - \lfloor y \rfloor, & \text{if } i = \lfloor y \rfloor + 1, \\ \lfloor y \rfloor + 1, & \text{if } i = \lfloor y \rfloor, \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Finally, we trained a logistic regression model that predicts the transformed target  $d$  from the sentence pair representations of the training dataset. The results for the STS Benchmark test dataset are summarized in Table 6. Figure 4 presents a scatter plot of the results for the proposed models and the target  $y$ .

The experimental results show that the P-thought models of all three levels outperformed the bench-



Table 6. Experimental results for the STS Benchmark task

Model	Pearson correlation
PV-DBOW (Le & Mikolov, 2014) (Lau & Baldwin, 2016)	0.649
SipThought (Kiros et al., 2015)	0.721
SIF (Arora et al., 2017)	0.720
Sent2vec (Pagliardini et al., 2017)	0.755
InferSent (Conneau et al., 2017a)	0.758
P-thought (one layer-Bi RNN)	<b>0.812</b>
P-thought (two layers-Forward RNN)	<b>0.797</b>
P-thought (two layers-Bi RNN)	<b>0.764</b>

marked models. An interesting observation is that the performances of the P-thought models for the STS benchmark task are inversely proportional to the model complexity: the simplest model (one-layer Bi RNN) yielded the highest correlation value, while the most complex model (two-layers Bi RNN) resulted in the lowest correlation value among the three P-thought models. This observation is exactly the opposite of the result for the MS-COCO dataset, where the more complex the P-thought model, the higher is the P-coherence score. One possible reason for this reversed performances is that the MS-COCO caption dataset used for the model training only contains around 600,000 sentences, which is far fewer than training datasets for general sequence learning tasks in the NLP field. Hence, it is more likely to overfit the training dataset for a more complex structure. This problem can be alleviated by obtaining more of paraphrase sentence pairs.

## 6. Conclusion

Sentence embedding is one of the most important text processing techniques in NLP. To date, various sentence embedding models have been proposed and have yielded good performances in document classification and sentiment analysis tasks. However, the fundamental ability of sentence embedding methods, i.e., how effectively the meanings of the original sentences are preserved in the embedded vectors, cannot be fully evaluated through such indirect methods.

In this study, under the proposition that a good sentence embedding method should act similar to human language recognition, we suggested the concept of se-

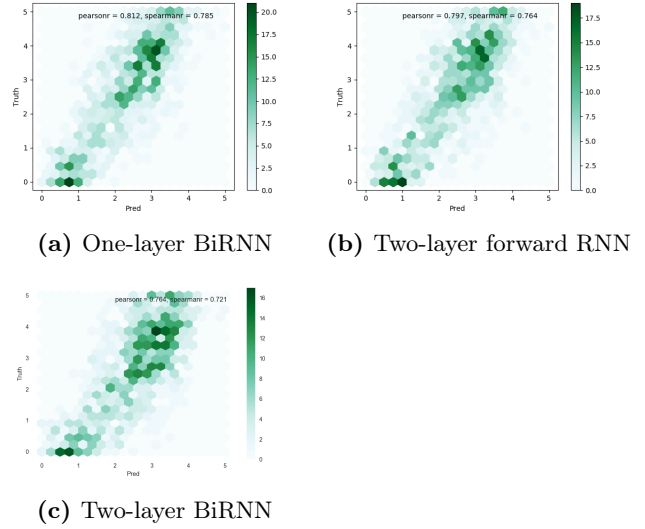


Figure 4. Scatter plot for the STS Benchmark task result

mantic coherence and proposed a model named P-thought that aims to maximize the semantic coherence by designing a model to have a dual generation structure. The proposed model was evaluated based on the MS-COCO caption and STS Benchmark datasets. Experimental results showed that the P-thought models yielded better performances than the benchmarked models for both tasks. Based on the scatter plots in the two-dimensional space reduced by  $t$ -SNE, it can clearly be observed that the paraphrase sentences are more concentrated for the P-thought models than those using other sentence embedding methods.

The main limitation of the current work is that there are insufficient paraphrase sentences for training the models. P-thought models with more complex encoder structures tend to overfit the MS-COCO datasets. Although this problem can be resolved by acquiring more paraphrase sentences, it is not easy in practice to obtain a large number of paraphrase sentences. Therefore, similar to the approaches that have achieved good performances in machine translation by employing semi-supervised learning or unsupervised learning (Cheng et al., 2016; Artetxe et al., 2017; Lample et al., 2017), an approach to improve the performances of the proposed models using only minimal paraphrase data should be developed.

## References

- Arora, Sanjeev, Liang, Yingyu, and Ma, Tengyu. A simple but tough-to-beat baseline for sentence embeddings. *International conference on Learning Representations*, 2017.



- Artetxe, Mikel, Labaka, Gorka, Agirre, Eneko, and Cho, Kyunghyun. Unsupervised neural machine translation. *arXiv preprint arXiv:1710.11041*, 2017.
- Bowman, Samuel R, Angeli, Gabor, Potts, Christopher, and Manning, Christopher D. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*, 2015.
- Cer, Daniel, Diab, Mona, Agirre, Eneko, Lopez-Gazpio, Inigo, and Specia, Lucia. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint arXiv:1708.00055*, 2017.
- Chen, Minmin. Efficient vector representation for documents through corruption. *arXiv preprint arXiv:1707.02377*, 2017.
- Cheng, Yong, Xu, Wei, He, Zhongjun, He, Wei, Wu, Hua, Sun, Maosong, and Liu, Yang. Semi-supervised learning for neural machine translation. *arXiv preprint arXiv:1606.04596*, 2016.
- 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. *arXiv preprint arXiv:1406.1078*, 2014.
- Conneau, Alexis, Kiela, Douwe, Schwenk, Holger, Barrault, Loic, and Bordes, Antoine. Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*, 2017a.
- Conneau, Alexis, Schwenk, Holger, Barrault, Loïc, and Lecun, Yann. Very deep convolutional networks for text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pp. 1107–1116, 2017b.
- Deng, Jia, Dong, Wei, Socher, Richard, Li, Li-Jia, Li, Kai, and Fei-Fei, Li. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255. IEEE, 2009.
- Glorot, Xavier and Bengio, Yoshua. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 249–256, 2010.
- Gupta, Ankush, Agarwal, Arvind, Singh, Prawaan, and Rai, Piyush. A deep generative framework for paraphrase generation. *arXiv preprint arXiv:1709.05074*, 2017.
- Hill, Felix, Cho, Kyunghyun, and Korhonen, Anna. Learning distributed representations of sentences from unlabelled data. *arXiv preprint arXiv:1602.03483*, 2016.
- Huang, Eric H, Socher, Richard, Manning, Christopher D, and Ng, Andrew Y. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pp. 873–882. Association for Computational Linguistics, 2012.
- Iyyer, Mohit, Manjunatha, Varun, Boyd-Graber, Jordan, and Daumé III, Hal. Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pp. 1681–1691, 2015.
- Karpathy, Andrej and Fei-Fei, Li. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3128–3137, 2015.
- Karpathy, Andrej, Joulin, Armand, and Fei-Fei, Li F. Deep fragment embeddings for bidirectional image sentence mapping. In *Advances in neural information processing systems*, pp. 1889–1897, 2014.
- Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kiros, Ryan, Zhu, Yukun, Salakhutdinov, Ruslan R, Zemel, Richard, Urtasun, Raquel, Torralba, Antonio, and Fidler, Sanja. Skip-thought vectors. In *Advances in neural information processing systems*, pp. 3294–3302, 2015.
- Lample, Guillaume, Denoyer, Ludovic, and Ranzato, Marc’Aurelio. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*, 2017.
- Lau, Jey Han and Baldwin, Timothy. An empirical evaluation of doc2vec with practical insights into document embedding generation. *arXiv preprint arXiv:1607.05368*, 2016.

- Le, Quoc and Mikolov, Tomas. Distributed representations of sentences and documents. In *International Conference on Machine Learning*, pp. 1188–1196, 2014.
- Lee, Jason, Cho, Kyunghyun, and Hofmann, Thomas. Fully character-level neural machine translation without explicit segmentation. *arXiv preprint arXiv:1610.03017*, 2016.
- Lin, Tsung-Yi, Maire, Michael, Belongie, Serge, Hays, James, Perona, Pietro, Ramanan, Deva, Dollár, Piotr, and Zitnick, C Lawrence. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Maaten, Laurens van der and Hinton, Geoffrey. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- Mikolov, Tomas, Le, Quoc V, and Sutskever, Ilya. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*, 2013a.
- Mikolov, Tomas, Sutskever, Ilya, Chen, Kai, Corrado, Greg S, and Dean, Jeff. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013b.
- Newman, David, Lau, Jey Han, Grieser, Karl, and Baldwin, Timothy. Automatic evaluation of topic coherence. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 100–108. Association for Computational Linguistics, 2010.
- Pagliardini, Matteo, Gupta, Prakhar, and Jaggi, Martin. Unsupervised learning of sentence embeddings using compositional n-gram features. *arXiv preprint arXiv:1703.02507*, 2017.
- Pennington, Jeffrey, Socher, Richard, and Manning, Christopher. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Prakash, Aaditya, Hasan, Sadid A, Lee, Kathy, Datla, Vivek, Qadir, Ashequl, Liu, Joey, and Farri, Oladimeji. Neural paraphrase generation with stacked residual lstm networks. *arXiv preprint arXiv:1610.03098*, 2016.
- Röder, Michael, Both, Andreas, and Hinneburg, Alexander. Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining*, pp. 399–408. ACM, 2015.
- Soumya George, K and Joseph, Shibily. Text classification by augmenting bag of words (bow) representation with co-occurrence feature. *IOSR Journal of Computer Engineering (IOSR-JCE) e-ISSN: 2278-0661, p-ISSN: 2278-8727*Volume, 16:34–38, 2014.
- Wan, Shengxian, Lan, Yanyan, Guo, Jiafeng, Xu, Jun, Pang, Liang, and Cheng, Xueqi. A deep architecture for semantic matching with multiple positional sentence representations. In *AAAI*, volume 16, pp. 2835–2841, 2016.
- Wieting, John, Bansal, Mohit, Gimpel, Kevin, and Livescu, Karen. Towards universal paraphrastic sentence embeddings. *arXiv preprint arXiv:1511.08198*, 2015.
- Xia, Yingce, Qin, Tao, Chen, Wei, Bian, Jiang, Yu, Nenghai, and Liu, Tie-Yan. Dual supervised learning. *arXiv preprint arXiv:1707.00415*, 2017.
- Zhao, Shenjian and Zhang, Zhihua. Deep character-level neural machine translation by learning morphology. 2016.
- Zhou, Peng, Qi, Zhenyu, Zheng, Suncong, Xu, Jiaming, Bao, Hongyun, and Xu, Bo. Text classification improved by integrating bidirectional lstm with two-dimensional max pooling. *arXiv preprint arXiv:1611.06639*, 2016.