

SEMANTIC SENTENCE EMBEDDINGS FOR PARAPHRASING AND TEXT SUMMARIZATION

Chi Zhang, Shagan Sah, Thang Nguyen, Dheeraj Peri, Alexander Loui[†], Carl Salvaggio, Raymond Ptucha

Rochester Institute of Technology, Rochester, NY 14623, USA

[†] Kodak Alaris Imaging Science R&D, Rochester, NY 14615, USA

ABSTRACT

This paper introduces a sentence to vector encoding framework suitable for advanced natural language processing. Our latent representation is shown to encode sentences with common semantic information with similar vector representations. The vector representation is extracted from an encoder-decoder model which is trained on sentence paraphrase pairs. We demonstrate the application of the sentence representations for two different tasks – sentence paraphrasing and paragraph summarization, making it attractive for commonly used recurrent frameworks that process text. Experimental results help gain insight how vector representations are suitable for advanced language embedding.

Index Terms— sentence embedding, sentence encoding, sentence paraphrasing, text summarization, deep learning

1. INTRODUCTION

Modeling temporal sequences of patterns requires the embedding of each pattern into a vector space. For example, by passing each frame of a video through a Convolutional Neural Network (CNN), a sequence of vectors can be obtained. These vectors are fed into a Recurrent Neural Network (RNN) to form a powerful descriptor for video annotation [1, 2, 3]. Similar, techniques such as word2vec [4] and GloVe [5] have been used to form vector representations of words. Using such embeddings, sentences become a sequence of word vectors. When these vector sequences are fed into a RNN, we get a powerful descriptor of a sentence [6].

Given vector representations of a sentence and video, the mapping between these vector spaces can be solved, forming a connection between visual and textual spaces. This enables tasks such as captioning, summarizing, and searching of images and video to become more intuitive for humans. By vectorizing paragraphs [7], similar methods can be used for richer textual descriptions.

Recent advances at vectorizing sentences represent exact sentences faithfully [8, 7, 9], or pair a current sentence with prior and next sentence [10]. Just like word2vec and GloVe map words of similar meaning close to one another, we desire a method to map sentences of similar meaning close to one

another. For example, the toy sentences “A man jumped over the stream” and “A person hurdled the creek” have similar meaning to humans, but are not close in traditional sentence vector representations. Just like the words flower, rose, and tulip are close in good sentence to vector representations, our toy sentences must lie close in the introduced embedded vector space.

Inspired by the METEOR [11] captioning benchmark which allows substitution of similar words, we choose to map similar sentences as close as possible. We utilize both paraphrase datasets and ground truth captions from multi-human captioning datasets. For example, the MS-COCO dataset [12] has over 120K images, each with five captions from five different evaluators. On average, each of these five captions from each image should convey the same semantic meaning. We present an encoder-decoder framework for sentence paraphrases, and generate the vector representation of sentences from this framework which maps sentences of similar semantic meaning nearby in the vector encoding space.

The main contributions of this paper are 1) The usage of sentences from widely available image and video captioning datasets to form sentence paraphrase pairs, whereby these pairs are used to train the encoder-decoder model; 2) We show the application of the sentence embeddings for paragraph summarization and sentence paraphrasing, whereby evaluations are performed using metrics, vector visualizations and qualitative human evaluation; and 3) We extend the vectorized sentence approach to a hierarchical architecture, enabling the encoding of more complex structures such as paragraphs for applications such as text summarization.

The rest of this paper is organized as follows: Section 2 reviews some related techniques. Section 3 presents the proposed encoder-decoder framework for sentence and paragraph paraphrasing. Section 4 discusses the experimental results. Concluding remarks are presented in Section 5.

2. RELATED WORK

Most machine learning algorithms require inputs to be represented by fixed-length feature vectors. This is a challenging task when the inputs are text sentences and paragraphs. Many studies have addressed this problem in both supervised and

unsupervised approaches. For example, [10] presented a sentence vector representation while [7] created a paragraph vector representation. An application of such representations is shown by [13], that has used individual sentence embeddings from a paragraph to search for relevant video segments.

An alternate approach uses an encoder-decoder [14] framework that first encodes f inputs, one at a time to the first layer of a two layer Long Short-Term Memory (LSTM), where f can be of variable length. Such an approach is shown for video captioning tasks by S2VT [15] that encodes the entire video, then decodes one word at a time.

There are numerous recent works on generating long textual paragraph summaries from videos. For example, [16] present a hierarchical recurrent network that comprise of a paragraph generator that is built on top of a sentence. The sentence generator encodes sentences into compact representations and the paragraph generator captures inter-sentence dependencies. [17] performed similar narratives from long videos by combining sentences using connective words at appropriate transitions learned using unsupervised learning.

3. METHODOLOGY

The vector representation of a sentence is extracted from an encoder-decoder model on sentence paraphrasing, and then tested on a text summarizer.

3.1. Vector Representation of Sentences

We consider the sentence paraphrasing framework as an encoder-decoder model, as shown in Fig. 1. Given a sentence, the encoder maps the sentence into a vector (sent2vec) and this vector is fed into the decoder to produce a paraphrase sentence. We represent the paraphrase sentence pairs as (S_x, S_y) . Let x_i denote the word embedding for sentence S_x ; and y_j denote the word embedding for sentence S_y , $i \in \{1 \dots T_x\}$, $j \in \{1 \dots T_y\}$ where T_x and T_y are the length of the paraphrase sentences.

Several choices for encoder have been explored, including LSTM, GRU [18] and BN-RHN [19]. In our model, we use an RNN encoder with LSTM cells since it is easy to be implemented and performs well on this model. Specifically, the words in S_x are converted into token IDs and then embedded using GloVe [5]. To encode a sentence, the embedded words are iteratively processed by the LSTM cell [14].

The decoder is a neural language model which conditions on the encoder output h_{T_x} . The computation is similar to that of the encoder. The vector h_{T_x} encodes the input sentence into a vector and is known as the vector representation of the input sentence, or “sent2vec” in this paper. Note that we do not have attention between encoder and decoder. This ensures that all the information extracted from the input sentence by encoder goes through sent2vec. In other words, attention is not adopted in order to avoid information leakage.




Fig. 1. The sentence paraphrasing model. The red and blue cells represent encoder and decoder respectively. The intermediate vector in black (sent2vec) is vector encoded sentence.

In full softmax training, for every training example in the word level, we would need to compute logits for all classes. However, this can get expensive if the universe of classes depending on size of vocabulary is very large. Given the predicted sentence and ground truth sentence, we use sampled softmax [20] as a candidate sampling algorithm. For each training sample, we pick a small set of sampled classes according to a chosen sampling function. A set of candidates C is created containing the union of the target class and the sampled classes. The training task figures out, given this set C , which of the classes in C is the target class.

3.2. Hierarchical Encoder for Text Summarization

Each sentence in a paragraph can be represented by a vector using the method described in Section 3.1. These vectors x_i , $i \in \{1 \dots T_x\}$ are fed into a hierarchical encoder [21], and then the summarized text is generated word by word in an RNN decoder, as shown in Fig. 2. We first divide all T_x vectors into several chunks (x_1, x_2, \dots, x_n) , $(x_{1+s}, x_{2+s}, \dots, x_{n+s})$, ..., $(x_{T-n+1}, x_{T-n+2}, \dots, x_T)$, where s is the stride and it denotes the number of temporal units adjacent chunks are apart. For each chunk, a feature vector is extracted using a LSTM layer and fed into the second layer. Each feature vector gives a proper abstract of its corresponding chunk. We also use LSTM units in the second layer to build the hierarchical encoder. The first LSTM layer serves as a filter and it is used to explore local temporal structure within subsequences. The second LSTM learns the temporal dependencies among subsequences. As a result, the feature vector generated from the second layer, which is called “paragraph2vec”, summarizes all input vectors extracted from the entire paragraph. Finally, a RNN decoder converts “paragraph2vec” into word sequence (y_1, y_2, \dots) , forming a summarized sentence.

We integrate a soft attention mechanism [22] in the hierarchical encoder. The attention mechanism allows the LSTM to pay attention to different temporal locations of the input sequence. When the input sequence and the output sequence are not strictly aligned, attention can be especially helpful.




Fig. 2. The paragraph summarizer. The red and blue cells represent encoder and decoder respectively. The encoder inputs x_i are the vector representation generated using sent2vec of the sentences in the paragraph. The decoder outputs y_i are the words in summarized text. The intermediate vector in black (paragraph2vec) is vector encoded paragraph.

4. EXPERIMENTAL RESULTS

4.1. Datasets

Visual Caption Datasets There are numerous datasets with multiple captions for images or videos. For example, MSR-VTT dataset [23] is comprised of 10,000 videos with 20 sentences each describing the videos. The 20 sentences are paraphrases since all the sentences are describing the same visual input. We form pairs of these sentences to create input-target samples. Likewise, MSVD [24], MS-COCO [25], and Flickr-30k [26] are used. Table 1 lists the statistics of datasets used. 5% of captions from all datasets was held out as a test set. In total, we created over 10M training samples.

Table 1. Sentence pairs statistics in captioning datasets.

	MSVD	MSRVTT	MSCOCO	Flickr
#sent	80K	200K	123K	158K
#sent/samp.	~42	20	5	5
# sent pairs	3.2 M	3.8 M	2.4 M	600 K

The SICK dataset We use the Sentences Involving Compositional Knowledge (SICK) dataset [27] as a test set for sentence paraphrasing task. It consists about 10,000 English sentence pairs, which are annotated for relatedness by means of crowd sourcing techniques. The sentence relatedness score (on a 5-point rating scale) for each sentence pair is provided and meant to quantify the degree of semantic relatedness between sentences.

TACoS Multi-Level Corpus We extract training pairs for the paragraph summarization task from TACoS Multi-Level Corpus [28]. This dataset provides coherent multi-sentence descriptions of complex videos featuring cooking activities with three levels of detail: “detailed”, “short” and “single sentence” description. There are 143 training and 42 test video sequences with ~ 20 annotations for each of the description

levels in each sequence.

4.2. Training Details

Sentence Paraphrasing We trained the model as described in Section 3.1 on the Visual Caption Datasets. The word embedding is initialized using GloVe [5]. The number of units per layer in both encoder and decoder are empirically set to 300 and 1024. We generate two sets of vocabularies with size $20k$ and $50k$. Stochastic gradient descent is employed for optimization, where the initial learning rate and decay factor are set to 0.0005 and 0.99, respectively.

Paragraph Summarization In this task we summarize “detailed” description to “single sentence” in TACoS Multi-Level Corpus. We select detailed descriptions with less than 20 sentences. There are total of 3,176 samples, 2,467 are used for training and 709 are used for testing. We employed the hierarchical architecture described in Section 3.2 with stride s of 5. 20 feature vectors (short paragraphs are zero padded) are fed into the model, with each vector’s sentence representation extracted from our paraphrasing model. To make the model more robust, soft attention is used between each layer. During training, we use learning rate 0.0001 and Adam optimizer. All the LSTM cells are set to 1024 units, except the one in sentence generation layer which is set to 256 units.

4.3. Sentence Paraphrasing

Given a reference sentence, the objective is to produce a semantically related sentence. The paraphrasing model was trained on Visual Caption Datasets and evaluated on the SICK dataset, without any fine-tuning. The results are shown in Table 2. The evaluation metrics for this experiment are Pearson’s r , Spearman’s ρ and Mean Squared Error(MSE). We use the same setup used by [29] for calculation of these metrics.

Table 2. Test set results on the SICK semantic relatedness task, where 300, 1024 denote the number of hidden units and 20k, 50k denote the size of the vocabulary. r and ρ are Pearson’s and Spearman’s metric respectively.

	r	ρ	MSE
sent2vec(300,20k)	0.7238	0.5707	0.4862
sent2vec(300,50k)	0.7472	0.5892	0.4520
sent2vec(1024,50k)	0.6673	0.5285	0.5679

In order to visualize the performance of our method, we applied PCA to the vector representation. Fig. 4 visualizes some of the paraphrase sentence pairs in the SICK dataset. Representations are sensitive to the semantic information of the sentences since pairwise sentences are close to each other. For example, point 2A and 4A are close because “watching” and “looking” are semantically related.

The semantic relatedness and grammar correctness are verified by human generated scores. Each score is the average of 32 different human annotators. Scores take values

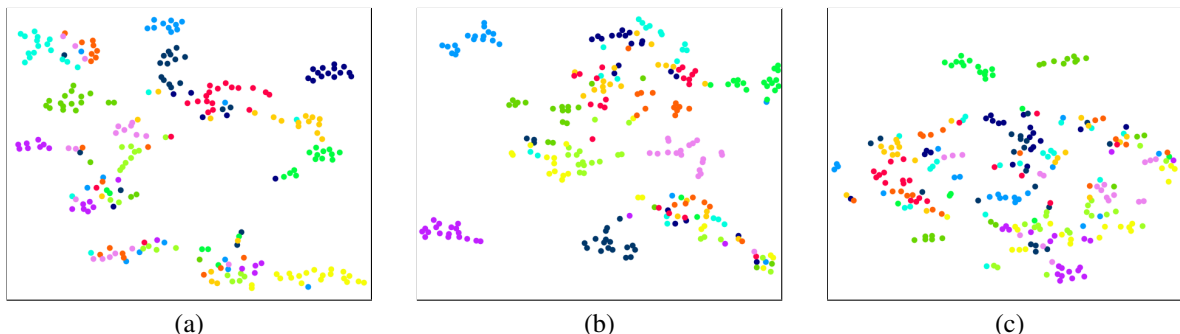


Fig. 3. t-SNE visualizations of single sentence descriptions of a subset of all sequences on TACoS Multi-Level Corpus. (a) Our sent2vec; (b) Skip-thoughts and (c) Skip-gram. Points are colored based on their sequence IDs. There are ~ 20 different annotations for each sequence. (Best viewed in color.)

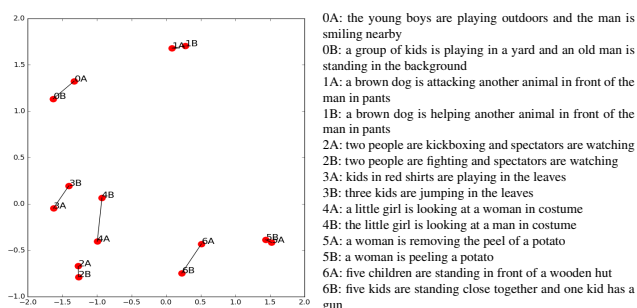


Fig. 4. Some paraphrase sentence pairs are represented by our sent2vec and then projected into 2D space using PCA. Each point represents a sentence in SICK dataset and the corresponding sentence is shown on the right.

between 1 and 5. A score of 1 indicates that the sentence pair is not at all related or totally incorrect syntax, while a score of 5 indicates they are highly related or grammatically correct. The sentences in human evaluation come from the Visual Caption and SICK test sets. The human evaluated scores of most sentence pairs are inversely proportional to the Euclidean distance between the vector representation of the corresponding sentences.

4.4. Text Summarization

We now show that in addition to paraphrasing, sent2vec is useful for text summarization. We use the TACoS Multi-Level Corpus [28]. Sentences from detailed descriptions of each video sequence are first converted into vectors using our model. These vectors are fed into our summarizer described in Section 3.2. The performance of the summarized text is evaluated based on the metric scores and compared to skip-thoughts and skip-gram. Note that skip-gram is used as the frequency-based average of word2vec for each word in the sentence. As shown in Table 3, the scores generated by our model are very close and comparable to the benchmark skip-thoughts. This result is reasonable since the dataset used in

training sent2vec are all from captions. The styles and topics of the sentences in this dataset are limited. However, the approach of forming sentence paraphrasing pairs and representing sentences using vectors are valid.

Table 3. Evaluation of short to single sentence summarization on TACoS Multi-Level Corpus using vectors from sent2vec, skip-thoughts, and skip-gram respectively.

	sent2vec	skip-gram	skip-thoughts
BLEU-1	0.479	0.514	0.520
BLEU-2	0.342	0.378	0.392
BLEU-3	0.213	0.245	0.276
BLEU-4	0.144	0.173	0.206
METEOR	0.237	0.250	0.253
ROUGE-L	0.48	0.509	0.522
CIDEr	1.129	1.430	1.562

Fig. 3 shows t-SNE [30] vector representation using (a) our sent2vec, (b) skip-thoughts and (c) skip-gram of randomly selected ~ 15 test sequences. In these plots, each point represents a single sentence. Points describing the same video sequence should be clustered. Points with the same color are nicely grouped in sent2vec visualization.

5. CONCLUSION

We showed the use of a deep LSTM based model in a sequence learning problem to encode sentences with common semantic information to similar vector representations. The presented latent representation of sentences has been shown useful for sentence paraphrasing and document summarization. We believe that reversing the encoder sentences helped the model learn long dependencies over long sentences. One of the advantages of our simple and straightforward representation is the applicability into a variety of tasks. Further research in this area can lead into higher quality vector representations that can be used for more challenging sequence learning tasks.

6. REFERENCES

- [1] Subhashini Venugopalan et al., “Translating videos to natural language using deep recurrent neural networks,” 2015.
- [2] Li Yao et al., “Describing videos by exploiting temporal structure,” in *ICCV*, 2015, pp. 4507–4515.
- [3] Subhashini Venugopalan et al., “Sequence to sequence – video to text,” in *ICCV*, 2015.
- [4] Tomas Mikolov et al., “Distributed representations of words and phrases and their compositionality,” in *NIPS*, 2013, pp. 3111–3119.
- [5] Jeffrey Pennington, Richard Socher, and Christopher D. Manning, “Glove: Global vectors for word representation,” in *EMNLP*, 2014, pp. 1532–1543.
- [6] Kyunghyun Cho et al., “On the properties of neural machine translation: Encoder-decoder approaches,” *arXiv preprint arXiv:1409.1259*, 2014.
- [7] Quoc V Le and Tomas Mikolov, “Distributed representations of sentences and documents.,” in *ICML*, 2014, vol. 14, pp. 1188–1196.
- [8] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom, “A convolutional neural network for modelling sentences,” *arXiv preprint arXiv:1404.2188*, 2014.
- [9] Han Zhao, Zhengdong Lu, and Pascal Poupart, “Self-adaptive hierarchical sentence model,” *arXiv preprint arXiv:1504.05070*, 2015.
- [10] Ryan Kiros et al., “Skip-thought vectors,” in *NIPS*, 2015, pp. 3294–3302.
- [11] Satanjeev Banerjee and Alon Lavie, “Meteor: An automatic metric for mt evaluation with improved correlation with human judgments,” in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005.
- [12] Tsung-Yi Lin et al., “Microsoft coco: Common objects in context,” in *ECCV*, 2014, pp. 740–755.
- [13] Jinsoo Choi et al., “Textually customized video summaries,” *arXiv preprint arXiv:1702.01528*, 2017.
- [14] Sutskever et al., “Sequence to sequence learning with neural networks,” in *NIPS*, 2014, pp. 3104–3112.
- [15] Subhashini Venugopalan et al., “Sequence to sequence-video to text,” in *ICCV*, 2015, pp. 4534–4542.
- [16] Haonan Yu et al., “Video paragraph captioning using hierarchical recurrent neural networks,” in *CVPR*, 2016, pp. 4584–4593.
- [17] Andrew Shin et al., “Beyond caption to narrative: Video captioning with multiple sentences,” in *ICIP*. IEEE, 2016, pp. 3364–3368.
- [18] Junyoung Chung et al., “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *arXiv preprint arXiv:1412.3555*, 2014.
- [19] Chi Zhang et al., “Batch normalized recurrent highway networks,” in *ICIP*. IEEE, 2017.
- [20] Sébastien Jean et al., “On using very large target vocabulary for neural machine translation,” *arXiv preprint arXiv:1412.2007*, 2014.
- [21] Pingbo Pan et al., “Hierarchical recurrent neural encoder for video representation with application to captioning,” in *CVPR*, 2016, pp. 1029–1038.
- [22] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [23] Jun Xu et al., “Msr-vtt: A large video description dataset for bridging video and language,” in *CVPR*, 2016.
- [24] David L Chen and William B Dolan, “Collecting highly parallel data for paraphrase evaluation,” in *Proceedings of the 49th Association for Computational Linguistics: Human Language Technologies-Volume 1*. ACL, 2011, pp. 190–200.
- [25] Tsung-Yi Lin et al., “Microsoft coco: Common objects in context,” in *European Conference on Computer Vision*. Springer, 2014, pp. 740–755.
- [26] Peter Young et al., “From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions,” *Transactions of the Association for Computational Linguistics*, 2014.
- [27] Marco Marelli et al., “A sick cure for the evaluation of compositional distributional semantic models.,” in *LREC*, 2014, pp. 216–223.
- [28] Anna Rohrbach et al., “Coherent multi-sentence video description with variable level of detail,” in *German Conference for Pattern Recognition*, 2014.
- [29] Christopher D. Manning Kai Sheng Tai, Richard Socher, “Improved semantic representations from tree-structured long short-term memory networks,” in *Association for Computational Linguistics*, 2015.
- [30] Laurens van der Maaten and Geoffrey Hinton, “Visualizing data using t-sne,” *Journal of Machine Learning Research*, vol. 9, no. Nov, pp. 2579–2605, 2008.