

CON-S2V: A Generic Framework for Incorporating Extra-Sentential Context into Sen2Vec

Tanay Kumar Saha¹ (orcid.org/0000-0002-7478-9222), Shafiq Joty²
(orcid.org/0000-0002-9222-2641), and Mohammad Al Hasan¹
(orcid.org/0000-0002-8279-1023)

¹ Indiana University Purdue University Indianapolis, Indianapolis, IN 46202, USA,
{tksaha, alhasan}@iupui.edu

² Nanyang Technological University, Singapore
rjoty@ntu.edu.sg

Abstract. We present a novel approach to learn distributed representation of sentences from unlabeled data by modeling both content and context of a sentence. The content model learns sentence representation by predicting its words. On the other hand, the context model comprises a neighbor prediction component and a regularizer to model distributional and proximity hypotheses, respectively. We propose an online algorithm to train the model components jointly. We evaluate the models in a setup, where contextual information is available. The experimental results on tasks involving classification, clustering, and ranking of sentences show that our model outperforms the best existing models by a wide margin across multiple datasets.

Keywords: Sen2Vec · Extra-Sentential Context · Embedding of Sentences

1 Introduction

For many text processing tasks that involve classification, clustering, or ranking of sentences, vector representation of sentences is a prerequisite. Bag-of-words (BOW) based vector representation has been used traditionally in these tasks, but in recent years, it has been shown that *distributed representation*, in the form of condensed real-valued vectors, learned from *unlabeled* data outperforms BOW based representations [?]. It is now well established that distributed representation captures semantic properties of linguistic units and yields better generalization [?,?].

However, most of the existing methods to devise distributed representation for sentences consider only the content of a sentence or its grammatical structure [?,?] disregarding its context. But, sentences rarely stand on their own in a text, rather the meaning of one sentence depends on the meaning of others within its context. For example, sentences in a text segment address a common topic [?]. At a finer level, sentences are connected by certain coherence relations (e.g., *elaboration*, *contrast*) and acts together to express a coherent message holistically [?].

Our work is built on the following hypothesis: *since the meaning of a sentence can be best interpreted within its context, its representation should also be inferred from its context*. Several recent works attempt to learn sentence representations which support

the above hypothesis by utilizing words or word sequences of neighboring sentences [?,?]. However, by learning representations to predict content of neighboring sentences, existing methods may learn semantic and syntactic properties that are more specific to the neighbors rather than the sentence under consideration. Furthermore, these methods either make a simple BOW assumption or disregard context when extracting a sentence vector.

In contrast to the existing works, we consider neighboring sentences as *atomic* linguistic units, and propose novel methods to learn the representations of a given sentence by jointly modeling content and context of a sentence. Our work considers two types of context: *discourse* and *similarity*. The discourse context of a given sentence \mathbf{v} comprises with its previous and the following sentence in the text. On the other hand, the similarity context is based on a user defined similarity function; thus it allows any sentences in the text to be in the context of \mathbf{v} depending on how similar that sentence is with \mathbf{v} based on the chosen function.

Our proposed computational model for learning the vector representation of a sentence comprises three components. The first component models the content by asking the sentence vector to predict its constituent words. The second component models the *distributional* hypotheses [?] of a context. The distributional hypothesis conveys that the sentences occurring in similar contexts should have similar representations. Our computation model captures this preference by using a context prediction component. Finally, the third component models the *proximity* hypotheses of a context, which also suggests that sentences that are proximal should have similar representations. Our method achieves this preference by using a Laplacian regularizer. To this end, we consider the sentence representation learning problem as an optimization problem whose objective function is built with expressions from the above three components and we solve this optimization problem by using an efficient online algorithm.

1.1 Summary of Results

We evaluate our sentence representation for learning models on multiple information retrieval tasks: topic classification and clustering, and single-document summarization. Our evaluation on these tasks across multiple datasets shows impressive results for our model, which outperforms the best existing models by up to 7.7 F_1 -score in classification, 15.1 V -score in clustering, 3.2 ROUGE-1 score in summarization. We found that the discourse context performs better on topic classification and clustering tasks, while similarity context performs better on summarization. We make our code ³ and pre-processed dataset ⁴ publicly available.

2 Related Work

Extensive research has been conducted on learning distributed representation of linguistic units both in supervised (task-specific) and in unsupervised (task-agnostic) settings. In this paper, we focus on learning *sentence* representations from *unlabeled* data.

³ <https://github.com/tksaha/con-s2v/tree/jointlearning>

⁴ https://www.dropbox.com/sh/ruhsi3c0unn0nko/AAAgVnZpojvXx91oQ21WP_MYa?dl=0

Two log-linear models are proposed in [?] for learning representations of words: continuous bag-of-words (CBOW) and continuous skip-gram. CBOW learns word representations by predicting a word given its (intra-sentential) context. The skip-gram model on the other hand learns representation of a word by predicting the words in a context. [?] proposed C-PHRASE, an extension of CBOW, where the context is extracted from a syntactic parse of the sentence. Simple averaging or addition of word vectors to construct sentence vectors often works well [?], and serves as baselines in our experiments.

CBOW and skip-gram models are extended in [?] to sentences and documents by proposing distributed memory (DM) and distributed bag-of-words (DBOW) models. In these models, similar to words, a sentence is mapped to an unique id and its representation is learned using contexts of words in the sentence. DM predicts a word given a context and the sentence id, where DBOW predicts all words in a context independently given the sentence id. Since these models are agnostic to sentence structure, they are quite fast to train. However, they disregard extra-sentential context of a sentence.

Sequential denoising autoencoder (SDAE) and FastSent are proposed in [?] for modeling sentences. SDAE employs an encoder-decoder framework, similar to neural machine translation (NMT) [?], to denoise an original sentence (target) from its corrupted version (source). FastSent is an additive model to learn sentence representation from word vectors. Given a sentence as BOW, it predicts the words of its adjacent sentences. The auto-encode version of FastSent also predicts the words of the current sentence. SDAE composes sentence vectors sequentially, but it disregards context of the sentence. FastSent, on the other hand, is a BOW model that considers neighboring sentences.

Another context-sensitive model is Skip-Thought [?], which uses the NMT framework to predict adjacent sentences (target) given a sentence (source). Since the encoder and the decoder use recurrent layers to compose vectors sequentially, SDAE and Skip-Thought are very slow to train. Furthermore, by learning representations to predict content of neighboring sentences, these methods (FastSent and Skip-Thought) may learn linguistic properties that are more specific to the neighbors rather than the sentence under consideration.

By contrast, we encode a sentence by treating it as an atomic unit like word, and similar to DBOW, we predict the words to model its content. Similarly, context is considered in our model by treating neighboring sentences as atomic units. This abstraction makes our model quite fast to train.

3 The Model

We hypothesize that the representation of a sentence depends not only on its content words, but also on other sentences in its context. It will be convenient to present our learning model using graph.

Let $G = (V, E)$ be a graph, where $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{|V|}\}$ represents the set of sentences in our corpus, and edge $(\mathbf{v}_i, \mathbf{v}_j) \in E$ reflects some relation between sentences \mathbf{v}_i and \mathbf{v}_j . A sentence $\mathbf{v}_i \in V$ is a sequence of words $(v_i^1, v_i^2, \dots, v_i^M)$, each coming from a dictionary \mathcal{D} . We define $\mathcal{N}(\mathbf{v}_i)$ as the set of neighboring sentences of




Fig. 1. Two instances (see (b) and (c)) of our model for learning representation of sentence \mathbf{v}_2 within a context of two other sentences: \mathbf{v}_1 and \mathbf{v}_3 (see (a)). Directed and undirected edges indicate prediction loss and regularization loss, respectively, and dashed edges indicate that the node being predicted is randomly sampled. (Collected from: 20news-bydate-train/misc.forsale/74732. The central topic is “forsale”).

\mathbf{v}_i , which constitutes extra-sentential context for sentence \mathbf{v}_i . We formalize relation between sentences and context later in Section 3.3.

Let $\phi : V \rightarrow \mathbb{R}^d$ be the mapping function from sentences to their distributed representations, i.e., real-valued vectors of d dimensions. Equivalently, ϕ can be thought of as a look-up matrix of size $|V| \times d$, where $|V|$ is the total number of sentences. Our aim is to learn $\phi(\mathbf{v}_i)$ by incorporating information from two different sources: (i) the content of the sentence, $\mathbf{v}_i = (v_i^1, v_i^2, \dots, v_i^M)$; and (ii) the context of the sentence in the graph, i.e., $\mathcal{N}(\mathbf{v}_i)$. Let $\langle v_i \rangle_t^l = (v_i^{t-l}, \dots, v_i^t, \dots, v_i^{t+l})$ denote a window of $2l + 1$ words around the word v_i^t in sentence \mathbf{v}_i , and $C_i = |\mathcal{N}(\mathbf{v}_i)|$ denote the context size for sentence \mathbf{v}_i . We define our model as a combination of three different loss functions:

$$J(\phi) = \sum_{\mathbf{v}_i \in V} \sum_{\substack{v \in \langle v_i \rangle_t^l \\ j \sim \mathcal{U}(1, C_i)}} [\mathcal{L}_c(\mathbf{v}_i, v) + \mathcal{L}_g(\mathbf{v}_i, \mathbf{v}_j) + \mathcal{L}_r(\mathbf{v}_i, \mathcal{N}(\mathbf{v}_i))] \quad (1)$$

where loss $\mathcal{L}_c(\mathbf{v}_i, v)$ is used to model the content of a sentence \mathbf{v}_i , and other two loss functions are for modeling the context of the sentence. We define $\mathcal{L}_c(\mathbf{v}_i, v)$ as the cost for predicting the content word v using the sentence vector $\phi(\mathbf{v}_i)$ as input features. Similarly, $\mathcal{L}_g(\mathbf{v}_i, \mathbf{v}_j)$ is defined as the cost for predicting a neighboring node $\mathbf{v}_j \in \mathcal{N}(\mathbf{v}_i)$, again using the sentence vector $\phi(\mathbf{v}_i)$ as input. The third loss $\mathcal{L}_r(\mathbf{v}_i, \mathcal{N}(\mathbf{v}_i))$ is a graph smoothing regularizer defined over the context of \mathbf{v}_i , which encourages two proximal sentences to have similar representations.

To learn the representation of a sentence \mathbf{v}_i using Eq. 1, for each content word v in a window $\langle v_i \rangle_t^l$, we sample a neighboring node \mathbf{v}_j from $\mathcal{N}(\mathbf{v}_i)$, uniformly at random, with replacement. We use the sentence vector $\phi(\mathbf{v}_i)$ (under estimation) to predict v and \mathbf{v}_j , respectively. A regularization is performed to smooth the estimated vector with respect to the neighboring vectors. Fig. 1 shows instances of our model for learning the representation of sentence \mathbf{v}_2 within a context of two other sentences: \mathbf{v}_1 and \mathbf{v}_3 .

We can use the standard *softmax* function for the prediction tasks. Formally, the negative log probability of an item o (can be a content word or a neighboring node) given the sentence vector $\phi(\mathbf{v}_i)$ is

$$-\log p(o|\mathbf{v}_i) = -\mathbf{w}_o^T \phi(\mathbf{v}_i) + \log \sum_{o' \in \mathcal{O}} \exp(\mathbf{w}_{o'}^T \phi(\mathbf{v}_i)) \quad (2)$$

where \mathcal{O} is the set of all possible items (i.e., vocabulary of words or set of all nodes), and \mathbf{w} 's are the weight parameters. Optimization is typically performed using gradient-based online methods, such as stochastic gradient descend (SGD), where gradients are obtained via backpropagation.

Unfortunately, training could be impractically slow on large corpora due to summation over all items in \mathcal{O} (Eq. 2), which needs to be performed for every training instance (\mathbf{v}_i, o) . Several methods have been proposed to address this issue including *hierarchical softmax* [?], *noise contrastive estimation* [?], and *negative sampling* [?]. We use negative sampling, which samples negative examples to approximate the summation term. Specifically, for each training instance (\mathbf{v}_i, o) , we add S negative examples $\{(\mathbf{v}_i, o^s)\}_{s=1}^S$ by sampling o^s from a known noise distribution ψ (e.g., *unigram*, *uniform*). The negative log probability in Eq. 2 is then formulated as such to discriminate a *positive* instance o from a *negative* one o^s :

$$-\log \sigma(\mathbf{w}_o^T \phi(\mathbf{v}_i)) - \log \sum_{s=1}^S \mathbb{E}_{o^s \sim \psi} \sigma(-\mathbf{w}_{o^s}^T \phi(\mathbf{v}_i)) \quad (3)$$

where σ is the sigmoid function defined as $\sigma(x) = 1/(1 + e^{-x})$, and \mathbf{w} 's and $\phi(\mathbf{v}_i)$ are similarly defined as before. Negative sampling thus reduces the number of computations needed from $|\mathcal{O}|$ to $S + 1$, where S is a small number (5 – 10) compared to the vocabulary size $|\mathcal{O}|$ (26K – 139K).

In the following, we elaborate on our methods for modeling content and context of a sentence.

3.1 Modeling Content

Our approach for modeling content of a sentence is similar to the distributed bag-of-words (DBOW) model of [?]. Given an input sentence \mathbf{v}_i , we first map it to a unique vector $\phi(\mathbf{v}_i)$ by looking up the corresponding vector in the sentence embedding matrix ϕ . We then use $\phi(\mathbf{v}_i)$ to predict each word v sampled from a window of words in \mathbf{v}_i . Formally, the loss for modeling content using negative sampling is

$$\mathcal{L}_c(\mathbf{v}_i, v) = -\log \sigma(\mathbf{w}_v^T \phi(\mathbf{v}_i)) - \log \sum_{s=1}^S \mathbb{E}_{v^s \sim \psi_c} \sigma(-\mathbf{w}_{v^s}^T \phi(\mathbf{v}_i)) \quad (4)$$

where σ is the sigmoid function as defined before, \mathbf{w}_v and \mathbf{w}_{v^s} are the weight vectors associated with words v and v^s , respectively, and ψ_c is the noise distribution from which v^s is sampled. In our experiments, we use unigram distribution of words raised to the 3/4 power as our noise distribution, in accordance to [?].

By asking the same sentence vector (under estimation) to predict its words, the content model captures the overall semantics of the sentence. The model has $O(d \times (|V| + |\mathcal{D}|))$ parameters.

3.2 Modeling Context

Our content model above attempts to capture the overall meaning of a sentence by looking at its words. However, sentences in a text are not independent, rather the meaning of a sentence depends on its neighboring sentences. For instance, consider the second and the third sentences in Fig. 1(a). When the sentences are considered in isolation, one cannot understand what they are talking about (i.e., *monitor for sale*). This suggests, since meaning of a sentence can be best interpreted within its context, the representation of the sentence should also be inferred from its context. We distinguish between two types of contextual relations between sentences: (i) distributional similarity, and (ii) proximity. Each of these corresponds to a loss in our model (Eq. 1), as we describe them below.

Modeling Distributional Similarity: Our sentence-level distributional hypothesis [?] is that if two sentences share many neighbors in the graph, their representations should be similar. We formulate this in our model by asking the sentence vector to predict its neighboring nodes. More formally, the loss for predicting a neighboring node $\mathbf{v}_j \in \mathcal{N}(\mathbf{v}_i)$ using the sentence vector $\phi(\mathbf{v}_i)$ is

$$\begin{aligned} \mathcal{L}_g(\mathbf{v}_i, \mathbf{v}_j) = & -\log \sigma \left(\mathbf{w}_j^T \phi(\mathbf{v}_i) \right) \\ & - \log \sum_{s=1}^S \mathbb{E}_{j^s \sim \psi_g} \sigma \left(-\mathbf{w}_{j^s}^T \phi(\mathbf{v}_i) \right) \end{aligned} \quad (5)$$

where \mathbf{w}_j and \mathbf{w}_{j^s} are the weight vectors associated with nodes \mathbf{v}_j and \mathbf{v}_{j^s} , respectively, and ψ_g is the noise distribution over nodes from which \mathbf{v}_{j^s} is sampled. Similar to our content model, ψ_g is defined as unigram distribution of nodes raised to the 3/4 power. The unigram distribution is computed based on the occurrences of the nodes in the neighborhood sets, $\{\mathcal{N}(\mathbf{v}_i)\}_{i=1}^{|V|}$. This model has $O(d \times (|V| + |V|))$ parameters.

Modeling Proximity: According to our proximity hypothesis, sentences that are proximal in their contexts, should have similar representations. We use a Laplacian regularizer to model this. Formally, the regularization loss for modeling proximity for a sentence \mathbf{v}_i in its context $\mathcal{N}(\mathbf{v}_i)$ is

$$\mathcal{L}_r(\mathbf{v}_i, \mathcal{N}(\mathbf{v}_i)) = \frac{\lambda}{C_i} \sum_{\mathbf{v}_k \in \mathcal{N}(\mathbf{v}_i)} \|\phi(\mathbf{v}_i) - \phi(\mathbf{v}_k)\|^2 \quad (6)$$

where $C_i = |\mathcal{N}(\mathbf{v}_i)|$ as defined before, and λ is a hyper-parameter to control regularization strength.

Rather than including the Laplacian as a regularizer in the objective function, another option is to first learn sentence embeddings using other components of the model (e.g., first two loss functions in Equation 1), and then *retrofit* them using the Laplacian as a post-processing step. [?] adopted this approach to incorporate lexical semantics (e.g., synonymy, hypernymy) into word representations. We compare our approach with retrofitting in Section 5.

3.3 Context Types

In this section we characterize context of a sentence. We distinguish between two types of context: discourse context and similarity context.

Discourse Context: The discourse context of a sentence is formed by the previous and the following sentences in the text. As explained before, the order of the sentences carries important information. For example, adjacent sentences in a text are logically connected by certain coherence relations (e.g., *elaboration*, *contrast*) to express the meaning [?]. On a coarser level, sentences in a text segment (e.g., paragraph) address a common (sub)topic [?]. The discourse context thus captures both coherence and topic structures of a text.

Similarity Context: While the discourse context covers important discourse phenomena like *coherence* and *cohesion* [?], some applications might require a context type that is based on more direct measures of similarity, and considers relations between all possible sentences in a document and possibly across multiple documents. For example, graph-based methods for topic segmentation [?] and summarization [?] rely on complete graphs of sentences, where edge weights represent cosine similarity between sentences. In an empirical evaluation of data structures for representing discourse coherence, [?] advocate for a graph representation of discourse allowing non-adjacent connections.

Our similarity context allows any other sentence in the corpus to be in the context of a sentence depending on how similar they are. To measure the similarity, we first represent the sentences with vectors learned by Sen2Vec [?], then we measure the cosine between the vectors. We restrict the context size of a sentence for computational efficiency, while still ensuring that it is informative enough. We achieve this by imposing two kinds of constraints. First, we set thresholds for intra- and across-document connections: sentences in a document are connected only if their similarity value is above 0.5, and sentences across documents are connected only if their similarity is above 0.8. Second, we allow up to 20 most similar neighbors.

3.4 Training

Algorithm 1 illustrates the SGD-based algorithm to train our model. We first initialize the model parameters; the sentence vectors ϕ are initialized with small random numbers sampled from uniform distribution $\mathcal{U}(-0.5/d, 0.5/d)$, and the weight parameters w 's are initialized with zero. We then compute the noise distributions ψ_c and ψ_g for $\mathcal{L}_c(\mathbf{v}_i, v)$ and $\mathcal{L}_g(\mathbf{v}_i, \mathbf{v}_j)$ losses, respectively.

We iterate over the sentences in our corpus in each epoch of SGD, as we learn their representations. Specifically, to estimate the representation of a sentence, for each word token in the sentence, we take three gradient steps to account for the three loss functions in Eq. 1. By making the same number of gradient updates, the algorithm weights equally the contributions of content and context.

Algorithm 1: Training CON-S2V with SGD

Input : set of sentences V , graph $G = (V, E)$

Output: learned sentence vectors ϕ

1. Initialize model parameters: ϕ and \mathbf{w} 's;

2. Compute noise distributions: ψ_c and ψ_g

3. **repeat**

for each sentence $\mathbf{v}_i \in V$ **do**

for each content word $v \in \mathbf{v}_i$ **do**

 a) Generate a positive pair (\mathbf{v}_i, v) and S negative pairs $\{(\mathbf{v}_i, v^s)\}_{s=1}^S$ using ψ_c ;

 b) Take a gradient step for $\mathcal{L}_c(\mathbf{v}_i, v)$;

 c) Sample a neighboring node \mathbf{v}_j from $\mathcal{N}(\mathbf{v}_i)$;

 d) Generate a positive pair $(\mathbf{v}_i, \mathbf{v}_j)$ and S negative pairs $\{(\mathbf{v}_i, \mathbf{v}_j^s)\}_{s=1}^S$ using

ψ_g ;

 e) Take a gradient step for $\mathcal{L}_g(\mathbf{v}_i, \mathbf{v}_j)$;

 f) Take a gradient step for $\mathcal{L}_r(\mathbf{v}_i, \mathcal{N}(\mathbf{v}_i))$;

end

end

until convergence;

4 Evaluation Tasks

Different methods have been proposed to evaluate sentence representation models [?]. However, unlike most existing methods, our model learns representation of a sentence by exploiting contextual information in addition to the content.⁵ To be able to evaluate our models, we thus require corpora of annotated sentences with ordering and document boundaries preserved, i.e., documents with sentence-level annotations. To the best of our knowledge, no previous work has used or released such corpora for learning sentence representation. In this work, we automatically create large corpora of documents with sentence-level topic annotations, which are then used to evaluate our models on topic *classification* and *clustering* tasks. Additionally, we evaluate our models on a *ranking* task of generating *extractive* single-document summaries. In the interest of coherence, we present the summarization task, followed by topic classification and clustering.

4.1 Extractive Summarization

Extractive summarization is often considered as a ranking problem, where the goal is to select the most important sentences to form an abridged version of the source document(s) [?]. Unsupervised methods are the predominant paradigm for determining sentence importance. We use the popular graph-based algorithm LexRank [?]. The input to LexRank is a graph, where nodes represent sentences and edges represent cosine similarity between *vector representations* (learned by models) of the two corresponding

⁵ For this reason, we did not evaluate our models on tasks previously used to evaluate sentence representation models.

sentences. We run the PageRank [?] on the graph to compute importance of each sentence in the graph.⁶ The top-ranked sentences are extracted as the summary sentences.

Data: We use the benchmark datasets from DUC-2001 and DUC-2002, and evaluate our representation models on the official task of generating a 100-words summary for each document in the datasets.⁷ The sentence representations are learned independently a priori from the same source documents. Table 1 shows some basic statistics about the datasets. For each document, 2-3 short (≈ 100 words) human authored reference summaries are available, which we use as gold summaries for automatic evaluation.

Metric: We use the widely used automatic evaluation metric ROUGE [?]

to evaluate the system-generated summaries. ROUGE computes n -gram recall between a system-generated summary and a set of human-authored reference summaries. Among the variants, ROUGE-1 (i.e., $n = 1$) has been shown to correlate well with human judgments for short summaries [?]. Therefore, we only report ROUGE-1 in this paper.

Dataset	#Doc.	#Avg. Sen.	#Avg. Sum.
DUC 2001	486	40	2.17
DUC 2002	471	28	2.04

Table 1. Basic statistics about the DUC datasets

4.2 Topic Classification and Clustering

We evaluate our models by measuring how effective the learned vectors are when they are used as features for classifying or clustering the sentences into topics. Text categorization has now become a standard in evaluating cross-lingual word embeddings [?]. We use a MaxEnt classifier and a K-means++ [?] clustering algorithm for classification and clustering tasks, respectively.

Data: We use the standard text categorization corpora: *Reuters-21578* and *20-Newsgroups*. Reuters-21578 (henceforth Reuters) is a collection of 21, 578 news documents covering 672 topics.⁸ 20-Newsgroups (henceforth Newsgroups) is a collection of about 20, 000 news articles organized into 20 different topics.⁹ We used the standard train-test splits (*ModApte* split for Reuters) split, and selected documents only from the 8 most frequent topics in both datasets. The selected topics for Reuters dataset are: *acq*, *crude*, *earn*, *grain*, *interest*, *money-fx*, *ship*, and *trade*. The topics selected for Newsgroups dataset are: *sci.space*, *sci.med*, *talk.politics.guns*, *talk.politics.mideast*, *rec.autos*, *rec.sport.baseball*, *comp.graphics*, and *soc.religion.christian*.

Dataset	#Doc.	Total #sen.	Annot. #sen	Train #sen.	Test #sen.	#Class
<i>Reuters</i>	9,001	42,192	13,305	7,738	3,618	8
<i>Newsgroups</i>	7,781	95,809	22,374	10,594	9,075	8

Table 2. Statistics about Reuters and Newsgroups.

⁶ The dumping factor in the PageRank was set to 0.85.

⁷ <http://www-nlpir.nist.gov/projects/duc/guidelines>

⁸ <http://kdd.ics.uci.edu/databases/reuters21578/>

⁹ <http://qwone.com/~jason/20Newsgroups/>

Generating Sentence-level Topic Annotations: As mentioned above, both *NewsGroups* and *Reuters* datasets come with document-level topic annotations. However, we need sentence-level annotations for our evaluation. One option is to assume that all the sentences of a document share the same topic label as the document. However, this naive assumption induces a lot of noise. Although sentences in a document collectively address a common topic, not all sentences are directly linked to that topic, rather they play supporting roles. To minimize this noise, we employ our extractive summarizer introduced in Section 4.1 to select the top 20% sentences of each document as representatives of the document, and assign them the same topic label as the topic of the document. We used Sen2Vec [?] representation to compute cosine similarity between two sentences in LexRank. Table 2 shows statistics of the resulting datasets. Note that the sentence vectors are learned independently from an entire dataset (#Total Sen. column in Table 2).

Metrics: We report raw accuracy, macro-averaged F_1 -score, and Cohen’s κ for comparing classification performance. For clustering, we report V-measure [?] and adjusted mutual information or AMI [?]. We use all the annotated sentences (train+test in Table 2) for comparing clustering performance.

5 Experiments

In this section, we present our experiments — the models we compare, their settings, and the results.

5.1 Models Compared

We compare our representation learning model against several baselines and existing models. We also experiment with a number of variations of our proposed model considering which components of the model are active, types of context, and how we incorporate the context. For clarity, in our tables we show results divided into five evaluation groups:

(I) Existing Distributed Models: This group includes Sen2Vec [?], W2V-avg, C-PHRASE [?], FastSent [?], and Skip-Thought [?].

We used Mikolov’s implementation¹⁰ of **Sen2Vec**, which gave better results than gensim’s version when validated on the sentiment treebank [?]. Following the recommendation by [?], we concatenate the vectors learned by DM and DBOW models. The concatenated vectors also performed better on our tasks.

For **W2V-avg**, we obtain a sentence vector by averaging the word vectors learned by training a skip-gram Word2Vec [?] on our training set. Since code for **C-PHRASE** is not publicly available, we use pre-trained word vectors (of 300 dimensions) available from author’s webpage.¹¹ We first add the word vectors to obtain a sentence vector, then we normalize the vector with l_2 normalization. Normalized vectors performed better on our tasks than the ones obtained by simple addition.

¹⁰ <https://code.google.com/archive/p/word2vec/>

¹¹ <http://clic.cimec.unitn.it/composes/cphrase-vectors.html>

We use the auto-encode version of **FastSent** (FastSent+AE) since it considers both content and context of a sentence. For **Skip-Thought**, we use the pre-trained combine-skip model that concatenates the vectors encoded by uni- and bi-skip models.¹² The resultant vectors are of 4800 dimensions. The model was originally trained on the BookCorpus¹³ with a vocabulary size of 20K words, however, it uses publicly available CBOW Word2Vec vectors to expand the vocabulary size to 930, 911 words.

(II) *Non-distributed Model*: We use **Tf-Idf** as our non-distributed baseline, where a sentence is represented by tf*idf weighting of its words.

(III) *Retrofitted Models*: We compare our approach of modeling context with the retrofitting method of [?]. We first learn sentence vectors using the content model only (i.e., by turning off contextual components in Eq. 1). Then we retrofit these vectors with the graph Laplacian $\mathcal{L}_r(\mathbf{v}_i, \mathcal{N}(\mathbf{v}_i))$ to encourage the revised vectors to be similar to the vectors of neighboring sentences and also similar to their prior representations. We consider two types of graph contexts: discourse (RET-dis) and similarity (RET-sim).

(IV) *Regularized Models*: We compare with a variant of our model, where the loss to capture distributional similarity $\mathcal{L}_g(\mathbf{v}_i, \mathbf{v}_j)$ is turned off. This model considers the same information as the retrofitting model (i.e., content and proximity), but trains the vectors in a single step. Its comparison with our complete model will tell us how much distributional similarity contributes to the overall performance. We define regularizers on two types of contexts: discourse (REG-dis) and similarity (REG-sim).

(V) *Our Models*: We experiment with two variants of our combined model, CON-S2V: one with discourse context (CON-S2V-dis), and the other with similarity context (CON-S2V-sim).

Dataset	Task	Sen2Vec	FastSent	W2V-avg	REG-sim	REG-dis	CON-S2V-sim	CON-S2V-dis
		(win. size)			(win. size, reg. str.)		(win. size, reg. str.)	
Reuters	clas.	8	10	10	(8, 1.0)	(8, 1.0)	(8, 0.8)	(8, 1.0)
	clus.	12	8	12	(12, 0.3)	(12, 1.0)	(12, 0.8)	(12, 0.8)
Newsgroups	clas.	10	8	10	(10, 1.0)	(10, 1.0)	(10, 1.0)	(10, 1.0)
	clus.	12	12	12	(12, 1.0)	(12, 1.0)	(12, 0.8)	(10, 1.0)
DUC 2001	sum.	10	12	12	(10, 0.8)	(10, 0.5)	(10, 0.3)	(10, 0.3)
DUC 2002	sum.	8	8	10	(8, 0.8)	(8, 0.3)	(8, 0.3)	(8, 0.3)

Table 3. Optimal values of the hyper-parameters for different models on different tasks.

¹² <https://github.com/ryankiros/skip-thoughts>

¹³ <http://yknzhu.wixsite.com/mbweb>

	Topic Classification Results						Topic Clustering Results			
	Reuters			Newsgroups			Reuters		Newsgroups	
	F_1	Acc	κ	F_1	Acc	κ	V	AMI	V	AMI
Sen2Vec	83.25	83.91	79.37	79.38	79.47	76.16	42.74	40.00	35.30	34.74
W2V-avg	(+) 2.06	(+) 1.91	(+) 2.51	(-) 0.42	(-) 0.44	(-) 0.50	(-) 11.96	(-) 10.18	(-) 17.90	(-) 18.50
C-PHRASE	(-) 2.33	(-) 2.01	(-) 2.78	(-) 2.49	(-) 2.38	(-) 2.86	(-) 11.94	(-) 10.80	(-) 1.70	(-) 1.44
FastSent	(-) 0.37	(-) 0.29	(-) 0.41	(-) 12.23	(-) 12.17	(-) 14.21	(-) 15.54	(-) 13.06	(-) 34.40	(-) 34.16
Skip-Thought	(-) 19.13	(-) 15.61	(-) 21.8	(-) 13.79	(-) 13.47	(-) 15.76	(-) 29.94	(-) 28.00	(-) 27.50	(-) 27.04
Tf-Idf	(-) 3.51	(-) 2.68	(-) 3.85	(-) 9.95	(-) 9.72	(-) 11.55	(-) 21.34	(-) 20.14	(-) 29.20	(-) 30.60
RET-sim	(+) 0.92	(+) 1.28	(+) 1.65	(+) 2.00	(+) 1.97	(+) 2.27	(+) 3.72	(+) 3.34	(+) 5.22	(+) 5.70
RET-dis	(+) 1.66	(+) 1.79	(+) 2.30	(+) 5.00	(+) 4.91	(+) 5.71	(+) 4.56	(+) 4.12	(+) 6.28	(+) 6.76
REG-sim	(+) 2.53	(+) 2.53	(+) 3.28	(+) 3.31	(+) 3.29	(+) 3.81	(+) 4.76	(+) 4.40	(+) 12.78	(+) 12.18
REG-dis	(+) 2.52	(+) 2.43	(+) 3.17	(+) 5.41	(+) 5.34	(+) 6.20	(+) 7.40	(+) 6.82	(+) 12.54	(+) 12.44
CON-S2V-sim	(+) 3.83	(+) 3.55	(+) 4.62	(+) 4.52	(+) 4.50	(+) 5.21	(+) 14.98	(+) 14.38	(+) 13.68	(+) 13.56
CON-S2V-dis	(+) 4.29	(+) 4.04	(+) 5.22	(+) 7.68	(+) 7.56	(+) 8.80	(+) 9.30	(+) 8.36	(+) 15.10	(+) 15.20

Table 4. Performance of our models on topic classification and clustering tasks in comparison to Sen2Vec.

5.2 Model Settings

The representation dimensions were set to 300 in DM and DBOW models. The concatenation of the two vectors yields 600 dimensions for Sen2Vec. For a fair comparison, the dimensions in all other models that we train (except pre-trained C-PHRASE and Skip-Thought) were fixed to 600. All the prediction-based models were trained with SGD. Retrofitting was done using iterative method [?] with 20 iterations. The number of noise samples (S) in negative sampling was set to 5. We also used subsampling of frequent words [?], which together with negative sampling give significant speed-ups in training.

For each dataset described in Section 4, we randomly selected 20% documents from the training set to form a held-out validation set on which we tune the hyper-parameters. *Window size* (k) is a hyper-parameter that is common to all models. The regularized models have an additional hyper-parameter, *regularization strength* (λ). We tuned with $k \in \{8, 10, 12\}$ and $\lambda \in \{0.3, 0.6, 0.8, 1\}$, and we optimized F_1 for classification, AMI for clustering, and ROUGE-1 for summarization. Table 3 shows the hyper-parameters and their optimal values for each task. We evaluated our models on the test sets with these optimal values. We ran each experiment five times and take the average of the evaluation measures to avoid any randomness in results.

5.3 Classification and Clustering Results

Table 4 shows the results of the models on topic classification and clustering tasks, respectively. The scores are shown in comparison to Sen2Vec.

Unsurprisingly, Sen2Vec outperforms Tf-Idf representation (row 6) by a good margin on both tasks. It gets improvements of up to 11.6 points on classification, and up to 30.6 points on clustering. This is inline with the finding of [?], and demonstrates the benefits of using distributed representation over sparse BOW representations.

Simple averaging of Word2Vec vectors performs quite well for classification, especially, on Reuters, where it outperforms Sen2Vec by 1.9 to 2.5 points. [?] also reported

similar findings on five out of six datasets. However, averaging does not perform well on clustering, where the scores are 10.2 to 18.5 points below than Sen2Vec.

Simple addition-based composition of C-PHRASE word vectors performs poorly on both tasks – lower than Sen2Vec by 2 to 3 points on classification and by 1.4 to 11.9 points on clustering.

Unexpectedly, FastSent and Skip-Thought perform quite poorly on both tasks. Skip-Thought, in particular, has the worst performance on both tasks. These results contradict the claim made by [?] that skip-thought vectors are generic. To investigate if the poor results are due to shift of domains (*book* vs. *news*), we also trained Skip-Thought on our training corpora with vector size 600 and vocabulary size 30K. The performance was even worse. We hypothesize, this is due to our training set size, which may not be enough for the heavy model. Also, Skip-Thought does not perform any inference to extract the vector using a context – although the model was trained to generate neighboring sentences, context was ignored when the encoder was used to extract the sentence vector.

Regarding FastSent, although its classification performance on Reuters is comparable to Sen2Vec, it performs poorly on Newsgroups, where the measures are 12.2 to 14.3 points lower than Sen2Vec. The differences get bigger in clustering. The reason could be that FastSent does not learn sentence representations directly, rather it simply adds the word vectors. Note that FastSent was outperformed by Tf-Idf in all classification tasks in [?]. Since both Skip-Thought and FastSent learn representations by predicting contents of adjacent sentences, the learned vectors might capture linguistic properties that are more specific to the neighbors.

We also experimented with SAE and SDAE auto-encoders proposed in [?]. However, they performed poorly on our tasks (thus not shown in the table). For example, SAE gave accuracies of around 40% on reuters and 18% on newsgroups. This is similar to what [?] observed. They propose to use pretrained word embeddings to improve the results. We did not achieve significant gains by using pretrained embeddings on our tasks.

Interestingly, the retrofitting and regularized models improve over Sen2Vec on both tasks, showing gains of up to 6.2 points on classification and up to 12.8 points on clustering. The improvements in most cases are significant. This demonstrates that proximity hypothesis is beneficial for these tasks.

When we compare regularized models with retrofitted ones, we observe that regularized models consistently outperform the retrofitted counterparts on both tasks with improvements of up to 1.6 points on classification and up to 7.6 points on clustering. This demonstrates that incorporating contextual information by means of regularization is more effective than retrofitting. This could be due to the fact that regularization approach induces contextual information while learning the vectors from scratch as opposed to revising them in a post-processing step.

Finally, we observe further improvements for our complete models (CON-S2V variants) on both tasks. Compared to the best regularized models, our models deliver improvements of up to 2.6 points on classification and up to 7.6 points on clustering. This demonstrates that by including the neighbor prediction component to model distributional similarity, our model captures complementary contextual information to what is

captured by the regularized models. A comparison between the context types reveals that discourse context is more beneficial than similarity context in most cases, especially for classification. For clustering, similarity context gives better results in a few cases (e.g., on Reuters). Overall, our best model outperforms the best existing model by up to 8.8 and 15.20 points on classification and clustering tasks, respectively.

5.4 Summarization Results

Table 5 shows ROUGE-1 scores of our models on DUC datasets for the summary length of 100 words. W2V-avg performs well achieving comparable score to Sen2Vec on DUC’01 and 1.4 points improvement on DUC’02. C-PHRASE outperforms Sen2Vec by 2.5 and 1.7 points on DUC’01 and DUC’02, respectively. FastSent and Skip-Thought again perform disappointingly. Sen2Vec outperforms FastSent by 4.15 and 7.53 points on DUC’01 and DUC’02, respectively. Skip-Thought performs comparably to Sen2Vec on DUC’01, but gets worse on DUC’02.

Interestingly, Tf-Idf performs quite well on this task. It gives the top score on DUC’01 (i.e., 48.7 ROUGE-1), and an improvement of 1.5 points over Sen2Vec on DUC’02. These results suggest that existing distributed representation methods are inferior to traditional methods in modeling aspects that are necessary for measuring sentence importance.

Retrofitted models give mixed results and fail to get significant improvement over Sen2Vec. On the other hand, with similarity context, regularized model improves over Sen2Vec by 2 to 3 points. This again suggests that regularization is a better method to incorporate context proximity. By including the neighbor prediction component to incorporate distributional similarity, our combined model improves the scores further; it achieves the second best result on DUC’01, and becomes top-performer on DUC’02. It is not surprising that similarity context is more suitable than discourse context for this task. From a context of topically similar sentences, our model learns representations that capture linguistic aspects related to information centrality. Given that the existing models fail to beat the Tf-Idf baseline on this task, our results are rather encouraging.

	DUC’01	DUC’02
Sen2Vec	43.88	54.01
W2V-avg	(-) 0.62	(+) 1.44
C-PHRASE	(+) 2.52	(+) 1.68
FastSent	(-) 4.15	(-) 7.53
Skip-Thought	(+) 0.88	(-) 2.65
Tf-Idf	(+) 4.83	(+) 1.51
RET-sim	(-) 0.62	(+) 0.42
RET-dis	(+) 0.45	(-) 0.37
REG-sim	(+) 2.90	(+) 2.02
REG-dis	(-) 1.92	(-) 8.77
CON-S2V-sim	(+) 3.16	(+) 2.71
CON-S2V-dis	(+) 1.15	(-) 4.46

Table 5. ROUGE-1 scores of the models on DUC datasets in comparison with Sen2Vec.

6 Discussion and Future Directions

We have presented a novel model to learn distributed representation of sentences by considering content as well as context of a sentence. Our results on tasks involving

classifying, clustering and ranking sentences confirm that extra-sentential contextual information is crucial for modeling sentences, and this information is best captured by our model that comprises a neighbor-based prediction component and a regularization component to capture distributional similarity and contextual proximity, respectively.

One important property of our model is that it encodes a sentence directly, and it considers neighboring sentences as atomic units. Apart from the improvements that we achieve in various tasks, this property makes our model quite efficient to train compared to *compositional* methods like encoder-decoder models (e.g., SDAE, Skip-Thought) that compose a sentence vector from the word vectors. Encoder-decoder approaches attempt to capture the structure of a sentence, which could be beneficial to model long distance relations between words (e.g., negation in sentiment classification). It would be interesting to see how our model compares with compositional models on sentiment classification task. However, this would require creating a new dataset of comments with sentence-level sentiment annotations. We intend to create such datasets and evaluate the models in the future.

Acknowledgments

This research is partially supported by Mohammad Hasan’s NSF CAREER Award (IIS-1149851).

References

1. Le, Q.V., Mikolov, T.: Distributed representations of sentences and documents. In: ICML. Volume 14. (2014) 1188–1196
2. Mikolov, T., Yih, W.t., Zweig, G.: Linguistic regularities in continuous space word representations. In: Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Atlanta, Georgia, Association for Computational Linguistics (June 2013) 746–751
3. Bengio, Y., Ducharme, R., Vincent, P., Janvin, C.: A neural probabilistic language model. *J. Mach. Learn. Res.* **3** (March 2003)
4. Socher, R., Lin, C.C.Y., Ng, A.Y., Manning, C.D.: Parsing natural scenes and natural language with recursive neural networks. In: Proceedings of the 28th International Conference on Machine Learning (ICML). (2011) 129–136
5. Stede, M.: Discourse processing. Morgan & Claypool Publishers (2011)
6. Hobbs, J.R.: Coherence and coreference. *Cognitive science* **3**(1) (1979) 67–90
7. Kiros, R., Zhu, Y., Salakhutdinov, R., Zemel, R.S., Torralba, A., Urtasun, R., Fidler, S.: Skip-thought vectors. In: Proceedings of the 28th International Conference on Neural Information Processing Systems. NIPS’15, Montreal, Canada, MIT Press (2015) 3294–3302
8. Hill, F., Cho, K., Korhonen, A.: Learning distributed representations of sentences from unlabelled data. In: Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego, California, Association for Computational Linguistics (June 2016) 1367–1377
9. Harris, Z.: Distributional structure. *Word* **10** (1954) 146–162
10. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)

11. Pham, N.T., Kruszewski, G., Lazaridou, A., Baroni, M.: Jointly optimizing word representations for lexical and sentential tasks with the C-PHRASE model. In: ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers. (2015) 971–981
12. Mitchell, J., Lapata, M.: Composition in distributional models of semantics. *Cognitive Science* **34**(8) (2010) 1388–1439
13. Cho, K., van Merriënboer, B., Gülçehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder–decoder for statistical machine translation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics (October 2014) 1724–1734
14. Morin, F., Bengio, Y.: Hierarchical probabilistic neural network language model. In Cowell, R.G., Ghahramani, Z., eds.: Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics, Society for Artificial Intelligence and Statistics (2005) 246–252
15. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In Teh, Y., Titterton, M., eds.: Proc. Int. Conf. on Artificial Intelligence and Statistics (AISTATS). Volume 9 of JMLR W&CP. (2010) 297–304
16. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Proceedings of the 26th International Conference on Neural Information Processing Systems. NIPS’13 (2013) 3111–3119
17. Faruqi, M., Dodge, J., Jauhar, S.K., Dyer, C., Hovy, E., Smith, N.A.: Retrofitting word vectors to semantic lexicons. In: Proc. of NAACL. (2015)
18. Halliday, M., Hasan, R. In: Cohesion in English. Longman, London (1976)
19. Malioutov, I., Barzilay, R.: Minimum cut model for spoken lecture segmentation. In: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics. ACL-44, Sydney, Australia, Association for Computational Linguistics (2006) 25–32
20. Erkan, G., Radev, D.R.: Lexrank: Graph-based lexical centrality as salience in text summarization. *J. Artif. Int. Res.* **22**(1) (December 2004) 457–479
21. Wolf, F., Gibson, E.: Representing Discourse Coherence: A Corpus-Based Study. *Computational Linguistics* **31** (June 2005) 249–288
22. Nenkova, A., McKeown, K.: Automatic summarization. *Foundations and Trends in Information Retrieval* **5**(23) (2011) 103–233
23. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: bringing order to the web. (1999)
24. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out: Proceedings of the ACL-04 workshop. Volume 8., Barcelona, Spain (2004)
25. Hermann, K.M., Blunsom, P.: Multilingual models for compositional distributed semantics. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Baltimore, Maryland, Association for Computational Linguistics (June 2014) 58–68
26. Arthur, D., Vassilvitskii, S.: k-means++: The advantages of careful seeding. In: Proc. of the eighteenth annual ACM-SIAM symposium on Discrete algorithms. (2007) 1027–1035
27. Rosenberg, A., Hirschberg, J.: V-measure: A conditional entropy-based external cluster evaluation measure. In: EMNLP-CoNLL. Volume 7. (2007) 410–420
28. Vinh, N.X., Epps, J., Bailey, J.: Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.* **11** (December 2010) 2837–2854
29. Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C.D., Ng, A.Y., Potts, C.: Recursive deep models for semantic compositionality over a sentiment treebank. In: Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, Stroudsburg, PA, Association for Computational Linguistics (October 2013) 1631–1642