

Evaluation of Unsupervised Compositional Representations

Hanan Aldarmaki

The George Washington University
aldarmaki@gwu.edu

Mona Diab

The George Washington University
mtdiab@gwu.edu

Abstract

We evaluated various compositional models, from bag-of-words representations to compositional RNN-based models, on several extrinsic supervised and unsupervised evaluation benchmarks. Our results confirm that weighted vector averaging can outperform context-sensitive models in most benchmarks, but structural features encoded in RNN models can also be useful in certain classification tasks. We analyzed some of the evaluation datasets to identify the aspects of meaning they measure and the characteristics of the various models that explain their performance variance.

1 Introduction

Distributed semantic models for words encode latent features that reflect semantic aspects and correlations among words. The goal of compositional semantic models is to induce latent semantic representations that encode the meaning of phrases, sentences, and paragraphs of variable lengths. Some neural architectures such as convolutional (Kim, 2014) and recursive networks (Socher et al., 2013) handle variable-length input by identifying shift-invariant features suitable for the classification problem at hand, which makes it possible to skip composition and work directly with the entire space of individual word embeddings. While such models can achieve excellent performance in supervised classification tasks such as sentiment analysis, we are interested in generic unsupervised fixed-length representations for variable-length text sequences so as to efficiently preserve essential semantic content for later use in various supervised and unsupervised settings.

Binary bag-of-words are simple and effective representations that serve as a strong baseline in several classification benchmarks (Wang and Manning, 2012). However, they do not exploit the distributional relationships among different words, which limits their applicability and generalization when training data are scarce. Additive compositional functions, such as word vector sum or average, are more effective in semantic similarity tasks even when compared with tensor-based compositional functions (Milajevs et al., 2014) and can outperform more complex and better tuned models based on recurrent neural architectures on out-of-domain data (Wieting et al., 2015a). Yet, averaging also has several drawbacks: unlike binary representations, the individual word identities are lost, and some words that do not carry semantic significance may end up being more prominently represented than essential words. Furthermore, additive compositional models disregard sentence structure and word order, which can lead to loss of semantic nuance. To alleviate the first issue, the weights of various words can be adjusted using word frequency statistics (Riedel et al., 2017) or by inducing context-sensitive weights using recurrent neural networks (Wieting and Gimpel, 2017), both of which have been shown to outperform vector averaging. Context-sensitive feed-forward neural models like the paragraph vector (Le and Mikolov, 2014) potentially incorporate word order, yet the training objective may not be sufficient to model deeper structure. Sequence encoder-decoder models, on the other hand, can be trained with various sentence-level objectives, such as neural machine translation (NMT) (Sutskever et al., 2014), predicting surrounding

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

sentences (i.e. skip-thought) (Kiros et al., 2015), or reconstruction of the input using denoising auto-encoders (Hill et al., 2016). These sequential models have been evaluated and compared against other models and baselines on several supervised and unsupervised tasks in Hill et al. (2016). The denoising autoencoder model and skip-thought both performed well in supervised tasks, while the NMT model performed worse than the baselines. All three performed poorly in unsupervised settings.

To bridge some of the gaps in evaluation, we evaluated a subset of models with increasing complexity, from binary bag-of-words to RNNs, on various supervised and unsupervised settings. Our objective is to evaluate compositional models against strong baselines and identify the elements that lead to performance gains. We evaluated binary vs. distributed features, weighted vs. unweighted averaging, three different word embedding models, and four context-sensitive models that optimize different objectives: the paragraph vector, the gated recurrent averaging network (Wieting and Gimpel, 2017), skip-thought, and an LSTM encoder trained on labeled natural language inference data (inferSent) (Conneau et al., 2017). We also analyzed the intrinsic structures of the various models by visual inspection and k-means clustering to gain insights into structural differences that may explain the variance in performance.

2 Background: Unsupervised Compositional Models

2.1 Baselines

The simplest way of representing a sentence is a binary bag-of-words representation, where each word is a feature in the vector space. This results in large and sparse representations that only account for the existence of individual words within a sentence, yet they have been shown to be effective in various supervised classification tasks, especially in combination with n -grams and Naive Bayes (NB) features (Wang and Manning, 2012). Let \vec{x}_i be the binary representation of sentence i , and $y_i \in \{0, 1\}$ its label. The log-count ratio \vec{r} is calculated as

$$\vec{r} = \log \frac{\vec{p}/\|\vec{p}\|}{\vec{q}/\|\vec{q}\|} \quad (1)$$

Where $\vec{p} = 1 + \sum_{i:y_i=1} \vec{x}_i$ and $\vec{q} = 1 + \sum_{i:y_i=0} \vec{x}_i$ are the smoothed count vectors for each class (i.e. the number of samples in the class that include each feature). The feature vectors are then modified using the element-wise product $\vec{x}_i \circ \vec{r}$. NB features identify the most discriminative words for each task, so using them results in task-specific rather than general representations. However, given the relative efficiency of this model, we include it as a baseline for comparison.

2.2 Word Embeddings and Composition Functions

Representations of variable-length sentences and paragraphs can be constructed by averaging the embeddings of all words within a sentence. However, simple averaging may not be the best approach since not all words within a sentence are semantically relevant. The following methods can be used to adjust the weights of words according to their frequency, assuming that frequent words have lower semantic content:

tf-idf-weighted Average The *term frequency-inverse document frequency* statistic measures the importance of a word to a document. We treat each sentence as a document and calculate the *idf* weight for term t as follows:

$$idf_t = \log \frac{N}{1 + n_t} \quad (2)$$

where N is the total number of sentences and n_t the number of sentences in which the term appears. Terms that appear in more documents have lower *idf* weights.

sif-weighted Average The *smooth inverse frequency* (Riedel et al., 2017) is an alternative measure for discounting the weights of frequent words as follows:

$$sif_t = \frac{a}{a + p(t)} \quad (3)$$

where a is a smoothing parameter and $p(t)$ is the relative frequency of the term in the training corpus. In addition, as proposed in (Riedel et al., 2017), we subtract the projection of the vectors on the first principal component which corresponds to syntactic features associated with common words.

2.2.1 Word Embeddings

Random word projections: we generated a random vector drawn from the standard normal distribution for each word in the vocabulary. The vector sum of random word vectors is a low-dimensional projection of binary bag-of-words vectors. .

Continuous Bag of Words: CBOW is an efficient log-linear model for learning word embeddings using a feed-forward neural network classifier that predicts a word given the surrounding words within a fixed context window (Mikolov et al., 2013a). In Schnabel et al. (2015) word embedding evaluation, CBOW outperformed other word embeddings in word relatedness and analogy tasks.

Global Vectors: GloVe is a global log-bilinear regression model (Pennington et al., 2014) that produces word embeddings using weighted matrix factorization of word co-occurrence probabilities.

Subword Information Skip-gram `si-skip` learns representations for n -grams of various lengths, and words are represented as sums of n -gram representations (Bojanowski et al., 2017). The learning architecture is based on the continuous skip-gram model (Mikolov et al., 2013b), which is trained by maximizing the conditional probability of context words within a fixed window with negative sampling. The model exploits the morphological variations within a language to learn more reliable representations, particularly for rare morphological variants.

2.3 Neural Compositional Models

Several models have been proposed to overcome some of the weaknesses of bag-of-words and additive representations, such as lack of structure. We evaluated the following context-sensitive models:

The Paragraph Vector: `doc2vec` distributed memory model (Le and Mikolov, 2014) constructs representations for sentences and paragraphs using a neural feedforward network that maximizes the conditional probability of words within a paragraph given a context window and the paragraph embedding, which is shared for all contexts generated from the same paragraph. After learning word and paragraph embeddings for the training corpus, the model learns representations for new paragraphs by fixing the model parameters and updating the paragraph embeddings using backpropagation. This additional training at inference time considerably increases the time complexity of the model compared to all others in this study.

Gated Recurrent Averaging Network: GRAN has been recently introduced to combine the benefits of LSTM networks and averaging, where the weights are computed along with the word and sentence representations (Wieting and Gimpel, 2017). The model is trained using aligned sentences that are assumed to be paraphrases to maximize the similarity of their representations against negative examples. The intuition is to make the averaging operation context-sensitive, resulting in a more powerful construction than simple averaging where all words are equally important. The model was shown to outperform averaging and LSTM models in semantic relatedness tasks.

Skip-Thought: The `skip-th` model is a sequence encoder-decoder trained by projecting sentences into fixed-length vectors, which in turn are used as input to a decoder that is trained to reconstruct surrounding sentences (Kiros et al., 2015), where the encoder and decoder are RNNs with GRU activations (Chung et al., 2014). The model is trained with contiguous sentences extracted from a collection of novels. After training, the model’s vocabulary is expanded by learning a linear mapping from pre-trained CBOW word embeddings to the vector space of the `skip-th` word embeddings.

Natural Language Inference Encoder: In `inferSent` (Conneau et al., 2017), a bidirectional LSTM encoder with max-pooling is trained jointly with an inference classifier trained on the Stanford Natural Language Inference (SNLI) dataset (Bowman et al., 2015), which is a large manually-annotated dataset of English sentence pairs and their inference labels: {entailment, contradiction, neutral}.

Dataset	Pair-wise Similarity			Sentiment Analysis					Newsgroup				TREC
	STS	SICK	MSRP	CR	MPQA	RT-s	Subj	IMDB	REL	SPO	COM	POL	
Train	5,749	4,934	4,076	3,775	10,606	10,662	10,000	25k	1,078	1,604	1,694	1,310	5,452
Test	1,379	4,906	1,725	–	–	–	–	25k	–	–	–	–	500
pos ratio	–	–	0.66	0.64	0.31	0.50	0.50	0.50	0.58	0.51	0.51	0.52	–
l	12	10	23	21	3	21	24	262	82	82	82	92	10

Table 1: Dataset statistics. **Train**: number of samples in the training set. **Test**: number of samples in the test set, if applicable (CV is applied otherwise). **pos ratio**: ratio of positive samples in the test set (or total for datasets with no splits). l : average length of all samples.

3 Evaluation Datasets

To evaluate the text representations, we used them as features in extrinsic supervised and unsupervised tasks that reflect various semantic aspects, which can be grouped in three categories: pairwise-similarity, sentiment analysis, and categorization. A summary of the dataset statistics is in Table 1.¹

3.1 Pairwise Similarity

Semantic Textual Similarity: the STS benchmark dataset (Cer et al., 2017) includes a collection of English sentence pairs and human-annotated similarity scores that range from 0 (unrelated sentences) to 5 (paraphrases). The dataset includes training, development, and test sets. This task can be performed without supervision by calculating the cosine similarity between two sentence vectors. We also evaluated the models in a supervised settings using linear regression, where the input vector is a concatenation of the element-wise product $u.v$ and absolute difference $|u - v|$ of each pair $\langle u, v \rangle$.

Sentences Involving Compositional Knowledge: SICK dataset is a benchmark for evaluating compositional models (Marelli et al., 2014). We evaluated the models on the relatedness subtask, which is constructed in a similar manner as STS benchmark.

Paraphrase Detection: This is a binary classification task that involves the identification of paraphrases in similar sentence pairs using the Microsoft Research Paraphrase Corpus, MSRP (Dolan et al., 2004). We evaluated the models in two ways: calculating the cosine similarity between the sentence pairs and classifying them as paraphrases if the similarity is larger than a threshold tuned from the training set. The second approach is to learn a logistic regression classifier using a concatenation of $u.v$ and $|u - v|$.

3.2 Sentiment Analysis and Text Categorization

Sentiment Analysis: We used the following binary classification tasks: **CR** customer product reviews (Hu and Liu, 2004), **MPQA** opinion polarity subtask (Wiebe et al., 2005), **RT-s** short movie reviews (Pang and Lee, 2005), **Subj** subjectivity/objectivity classification task (Pang and Lee, 2004), and **IMDB** full-length movie review dataset (Maas et al., 2011).

Newsgroups: Following the setup in (Wang and Manning, 2012) we used the 20-Newsgroup dataset² to extract several binary topic categorization tasks. We processed the datasets to remove headers, forwarded text, and signatures, which results in smaller sentences and paragraphs. We used the following newsgroups for binary classification: **religion** (atheism vs. religion), **sports** (baseball vs. hockey), **computer** (windows vs. graphics), and **politics** (middle east vs. guns). We also trained multi-class classifiers on the 8 newsgroups.

Question Classification: We used the TREC 10 coarse question categorization task³ which categorizes questions into 6 classes: human (HUM), entity (ENTY), location (LOC), number (NUM), description (DESC), and abbreviation (ABBR).

¹Evaluation scripts and data can be downloaded from: https://github.com/h-aldarmaki/sentence_eval

²<http://qwone.com/~jason/20Newsgroups/>

³<http://cogcomp.org/Data/QA/QC/>

4 Experimental Setup

4.1 Training Data

We trained the unsupervised word embedding models CBOW, GloVe, and *si-skip* on a set of ~ 7 million sentences extracted from the English Wikipedia and Amazon movie and product reviews (He and McAuley, 2016). We also trained the Paragraph Vector (*doc2vec*) model on this dataset, and initialized the word embeddings using the *si-skip* pre-trained word embeddings above. While better results overall could be obtained using pre-trained word embeddings trained with much larger text corpora, we used this medium-size corpus to evaluate the various models consistently and reduce model variability due to data and vocabulary coverage.

We used the publicly available pre-trained GRAN⁴ and *skip-th*⁵ models, which require training with special types of datasets: paraphrase collections, and contiguous text from books, respectively. To ensure a fair evaluation, we only compared these models against binary bag-of-words and equivalent word embeddings. The word embeddings within the GRAN model were initialized with PARAGRAM-SL999 word vectors (Wieting et al., 2015b), so we used them as an evaluation baseline for GRAN. We compared *skip-th* against the CBOW embeddings that were used to expand the vocabulary, which account for most words in the final model’s vocabulary. We used the pre-trained *inferSent* model⁶ which uses pre-trained GloVe word embeddings⁷. We also experimented with the post-trained word embeddings for each model with similar results, so we omitted them for brevity.

4.2 Training Settings

We trained the unsupervised word embedding models using the optimal parameters recommended for each model. The hyper-parameters in *doc2vec* were set according to the recommendations in (Lau and Baldwin, 2016). For the supervised sentiment classification and text categorization tasks, we trained and tuned linear SVM models using grid search for datasets that include train/dev/test splits, and nested cross-validation otherwise. We also experimented with kernel SVMs but didn’t observe notable differences in the results.

5 Evaluation Results

5.1 Pairwise Similarity Evaluation

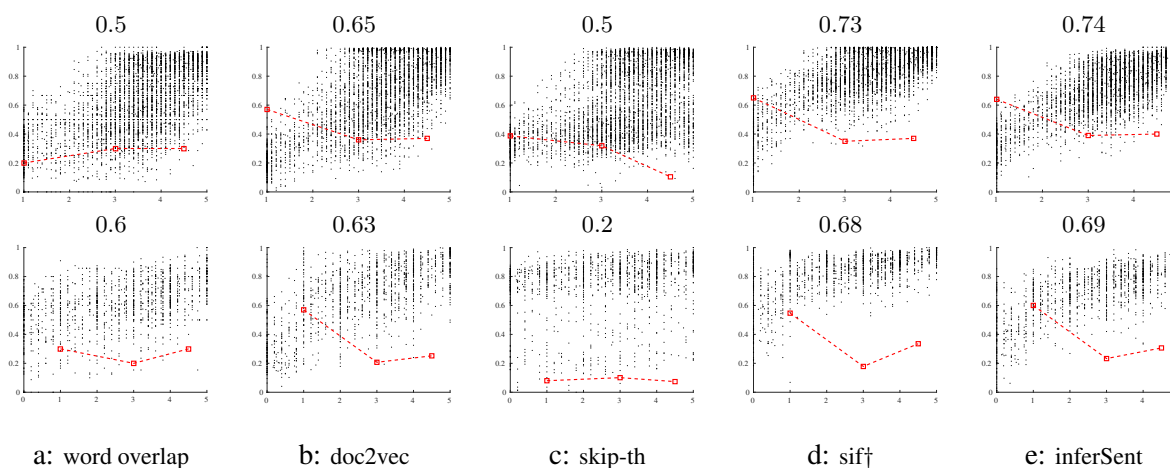


Figure 1: Scatter plots of normalized gold scores in the x axis vs. (a) word overlap (%) and (b - e) cosine similarity using various models. Top: SICK. Bottom: STS Benchmark. Pearson ρ plotted in red for $score \leq 2$, $score \in (2, 4)$, and $score \geq 4$. Overall Pearson ρ shown at the top. † sif-weighted average of pre-trained Glove vectors used in *inferSent*.

⁴<https://github.com/jwieting/acl2017>

⁵<https://github.com/ryankiros/skip-ths>

⁶<https://github.com/facebookresearch/InferSent>

⁷<https://nlp.stanford.edu/projects/glove>

		STS Benchmark ρ		SICK ρ		MSRP accuracy/F1	
		cosine	linear reg.	cosine	linear reg.	cosine	logistic reg.
Binary BOW		0.536	0.606	0.611	0.761	66.9/0.765	72.2/0.811
Paragraph Vector (doc2vec)		0.628	0.673	0.654	0.655	68.6/0.797	70.4/0.803
Random	avg	0.558	0.616	0.602	0.669	70.6/0.780	70.8/0.797
	idf	0.668	0.665	0.617	0.659	70.0/0.790	69.1/0.791
	sif	0.666	0.665	0.628	0.655	70.1/0.786	69.9/0.699
CBOW	avg	0.630	0.672	0.679	0.728	71.9/0.815	72.0/0.807
	idf	0.697	0.695	0.678	0.712	71.8/0.815	72.3/0.809
	sif	0.683	0.686	0.690	0.715	72.2/0.814	71.4/0.804
GloVe	avg	0.336	0.574	0.602	0.694	68.9/0.807	71.0/0.810
	idf	0.540	0.656	0.624	0.685	71.3/0.818	73.4/0.820
	sif	0.685	0.665	0.701	0.695	71.8/0.809	72.1/0.811
si-skip	avg	0.608	0.690	0.684	0.730	72.1/0.817	71.9/0.895
	idf	0.683	0.714	0.702	0.715	69.3/0.809	70.3/0.815
	sif	0.694	0.721	0.716	0.721	70.6/0.804	70.6/0.802
GRAN		0.747	0.747	0.715	0.756	71.3/0.817	72.3/0.812
Pre-trained PARAGRAM-SL999†	avg	0.564	0.690	0.694	0.746	71.8/0.818	73.2/0.816
	idf	0.711	0.733	0.723	0.756	72.1/0.817	73.3/0.817
	sif	0.716	0.722	0.733	0.765	73.4/0.822	72.0/0.809
skip-th		0.213	0.729	0.498	0.811	62.3/0.761	73.0/0.812
Pre-trained CBOW†	avg	0.631	0.695	0.727	0.758	70.3/0.813	73.2/0.813
	idf	0.674	0.708	0.710	0.731	69.6/0.809	71.3/0.807
	sif	0.686	0.707	0.727	0.737	69.8/0.809	70.9/0.803
inferSent		0.692	0.773	0.744	0.865	0.697/0.806	0.746/0.827
Pre-trained GloVe†	avg	0.497	0.655	0.687	0.753	0.711/0.818	0.732/0.817
	idf	0.606	0.688	0.696	0.736	0.688/0.809	0.711/0.804
	sif	0.679	0.699	0.729	0.749	0.709/0.816	0.708/0.804

Table 2: Pearson ρ for STS Benchmark and SICK relatedness, and Accuracy%/F1 for MSR Paraphrase detection. Results are shaded according to their statistical significance using the Williams test (Graham and Baldwin, 2014) with $\alpha = 0.05$. † pre-trained vectors used in the model above.

Table 2 shows the performance of the various models in the pair-wise similarity tasks. We highlight the best performance in each block; differences within the same shade are not statistically significant. Among the word embedding models, the subword skipgram `si-skip` achieved the best overall performance. For all models except `si-skip`, simple averaging performed poorly in semantic relatedness tasks, especially in the unsupervised setting, while `sif`-weighting generally outperformed `idf`-weighting (the improvement is most evident for `GloVe`). A similar trend is observed with random word vectors, which performed on par with `doc2vec`.

The binary bag-of-words model performed particularly well in the supervised SICK task. The performance of binary and random vectors can be explained by the high correlation between the percentage of overlapping words and the similarity scores as seen in Figure 1. We also highlighted the Pearson correlation coefficients for the following subsets of relatedness scores: $A = \{score \leq 2\}$, $B = \{score \in (2, 4)\}$, and $C = \{score \geq 4\}$. The overall Pearson correlation mostly reflects the performance on the most and least similar pairs, which tend to be the pairs with the highest and lowest word overlap, respectively; within the regions we highlighted, all correlations were relatively low. However, distributed models like `sif` and `inferSent` improved the correlation of A for SICK, and both A and C for STS Benchmark. Table 3 shows some examples of sentence pairs from A and C ; while both `sif` and `doc2vec` vectors consistently identified similar concepts (namely food related and competition concepts) regardless of surface similarity, binary scores only reflected the lexical similarity, which resulted in inconsistent scores.

Sentence 1	Sentence 2	Score	Binary	si-sif†	doc2vec	skip-th	infer	gl-sif†
A person is <u>frying</u> some food.	There is no person <u>peeling</u> a potato.	0.25	0.41	0.49	0.41	0.51	0.67	0.71
A woman is <u>cutting</u> a fish.	The main is <u>slicing</u> potatoes.	0.25	0.13	0.46	0.42	0.53	0.65	0.57
Two groups of people are <u>playing</u> football	Two team are <u>competing</u> in a football match.	0.93	0.52	0.74	0.72	0.60	0.79	0.71
Different teams are <u>playing</u> football on the field	Two teams are <u>playing</u> soccer	0.70	0.56	0.93	0.89	0.51	0.82	0.91

Table 3: Examples of sentence pairs in SICK and their relatedness scores vs. cosine similarity scores. All scores were normalized to be in $[0, 1]$. Shared words are shown in bold and related words underlined. † `sif`-weighted average of `si-skipgram` and pre-trained `GloVe`

Sentence 1	Sentence 2	Label
Bashir felt he was being tried by opinion not on the facts , Mahendradatta told Reuters.	Bashir also felt he was being tried by opinion rather than facts of law, he added.	1
West Nile Virus -which is spread through infected mosquitoes -is potentially fatal.	West Nile is a bird virus that is spread to people by mosquitoes .	0
SCO <u>says</u> the pricing terms for a license will not be announced for week	Details on pricing <u>will be announced</u> within a few weeks , <u>McBride said</u>	1
Russ Britt is the Los Angeles Bureau Chief for CBS.MarketWatch.com.	Emily Church is London bureau chief of CBS.MarketWatch.com.	0

Table 4: Examples of sentence pairs in MSRP and their labels. Shared words are shown in bold and related words underlined.

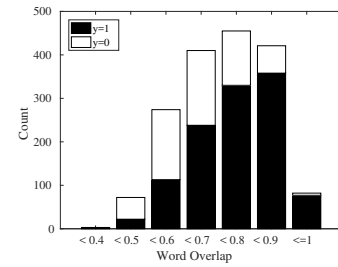


Figure 2: Ratio of paraphrases with increasing word overlap in MSRP.

GRAN outperformed simple averaging in both the STS and SICK tasks, which confirms the results in (Wieting and Gimpel, 2017), but compared with *idf* and *sif* averaging, there is no apparent improvement; it only outperformed weighted averaging in the unsupervised STS benchmark. *skip-th* vectors performed poorly in the unsupervised similarity tasks, but outperformed the pre-trained vectors in the supervised similarity tasks, particularly in SICK.

The low variance of the performance in the paraphrase detection task also reflects the overall correlation between word overlap and the likelihood of being a paraphrase as seen in Figure 2; for difficult cases, as in the examples in Table 4, the overall similarity is not a good indication of being a paraphrase. Significant improvements in this task may require more nuanced features as in (Ji and Eisenstein, 2013).

5.2 Evaluation on Sentiment Analysis and Categorization

		Sentiment Analysis					Text Categorization					
		CR	mpqa	RT-s	subj	imdb	rel	spo	com	pol	mult	TREC
Binary BOW		77.0/0.821	85.9/0.756	74.8	89.5	84.1	66.2/0.710	85.7	78.2	81.6	72.2	89.8
Unigram NBSVM		80.5	85.3	78.1	92.4	88.3	73.2	93.5	86.7	91.9	93.4	89.8
Paragraph Vector (doc2vec)		76.6/0.831	82.4/0.688	78.6	89.9	87.8	68.9/0.751	89.6	82.3	90.6	76.1	59.6
si-skip	avg	81.3/0.857	87.2/0.778	78.8	91.6	88.0	67.1/0.759	87.8	80.3	87.2	74.4	79.6
	idf	80.2/0.849	86.2/0.765	78.5	91.0	87.0	69.1/0.755	88.3	81.9	89.7	75.0	70.4
	sif	80.6/0.852	86.7/0.774	78.6	91.0	88.2	69.1/0.765	88.8	81.0	89.4	74.7	71.8
CBOW	avg	81.6/0.859	86.1/0.759	78.1	91.0	87.2	63.6/0.745	74.9	78.2	84.7	66.6	82.8
	idf	81.2/0.856	86.0/0.761	77.7	90.5	87.3	65.8/0.745	78.6	79.8	85.6	68.2	77.4
	sif	80.8/0.852	85.7/0.756	77.8	90.2	87.4	65.6/0.742	80.0	79.5	85.	68.3	76.2
GloVe	avg	80.7/0.851	85.4/0.745	77.6	91.0	87.3	67.1/0.748	82.1	78.8	86.5	69.3	79.8
	idf	80.7/0.851	85.8/0.756	77.8	90.8	87.5	65.5/0.725	85.2	77.7	87.6	70.9	72.4
	sif	80.6/0.850	85.4/0.751	77.6	90.7	87.4	66.8/0.736	85.8	78.5	87.5	70.9	72.0
Random	avg	70.7/0.779	74.2/0.413	61.9	77.3	74.0	55.8/0.634	71.4	72.8	69.4	47.	70.2
	idf	68.6/0.763	74.1/0.423	61.1	72.7	74.2	58.9/0.654	74.4	72.7	72.9	47.2	57.0
	sif	69.2/0.770	72.7/0.369	61.5	69.6	74.5	59.4/0.655	73.0	72.1	72.6	47.4	54.8
GRAN		78.4/0.838	86.6/0.769	75.1	88.5	83.1	66.0/0.753	90.6	80.8	88.5	73.2	60.4
pre-trained PARAGRAM-SL999†	avg	79.8/0.845	87.8/0.794	75.9	89.6	84.5	65.3/0.721	89.8	78.9	87.5	72.5	83.2
	idf	79.1/0.840	87.4/0.791	75.8	89.3	84.1	68.2/0.737	90.3	79.8	89.0	74.1	74.6
	sif	78.4/0.838	86.6/0.769	75.1	88.5	83.1	66.0/0.753	90.6	80.8	88.5	73.2	60.4
skip-th		80.4/0.851	87.0/0.78	76.4	93.4	81.8	65.5/0.736	70.4	69.4	81.5	60.1	88.2
Pre-trained CBOW †	avg	79.9/0.847	88.2/0.800	77.5	90.5	85.6	64.8/0.751	86.6	79.5	85.9	70.7	80.0
	idf	79.6/0.844	87.9/0.797	77.2	90.0	85.6	68.4/0.766	87.5	80.5	85.8	72.6	72.2
	sif	79.2/0.842	87.7/0.794	77.0	89.7	85.8	67.7/0.761	87.8	81.3	86.9	72.9	74.2
inferSent		83.0/0.867	88.5/0.811	77.1	91.0	86.4	68.5/0.731	88.6	80.9	85.2	74.4	88.6
Pre-trained GloVe †	avg	80.6/0.851	87.9/0.793	77.1	90.9	85.7	69.5/0.759	90.7	83.8	88.3	75.8	82.2
	idf	79.8/0.844	87.4/0.788	77.2	90.1	85.5	69.7/0.757	89.8	83.0	88.8	76.9	75.4
	sif	79.6/0.842	87.2/0.785	77.5	90.0	85.5	67.8/0.742	89.7	83.3	88.7	76.7	77.0

Table 5: Accuracy % or accuracy/F1 (for unbalanced datasets) on sentiment and topic categorization tasks. Results are shaded according to their statistical significance using a two-tailed significance test with $\alpha = 0.05$. † pre-trained word embeddings used in the model above

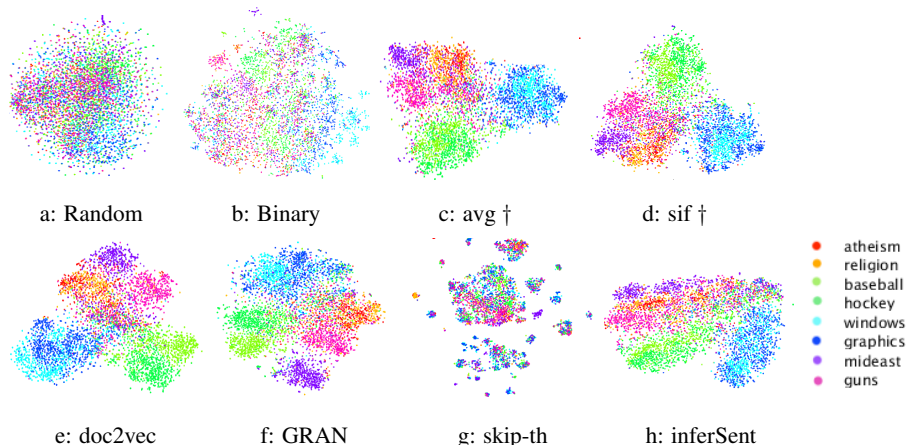


Figure 3: t-SNE visualizations of vectors on the 20-Newsdataset. † avg and sif using si-skip.

Table 5 shows the performance in sentiment analysis and categorization tasks. Unlike in pair-wise similarity, random vectors underperformed the NBSVM and distributed models by a large margin. This underscores the importance of global and distributed features in these tasks. si-skip outperformed other word embedding models, but we observe no advantage for weighted vs. unweighted averaging. In TREC question classification and the subjectivity benchmarks, avg performed significantly better than both idf and sif weighted averaging. skip-th vectors also significantly outperformed the pre-trained vectors in these two tasks, and underperformed in all others. We surmise that the syntactic features conveyed in frequent function words and the overall structure encoded by the LSTM network in skip-th may provide useful clues for these two classification tasks. inferSent achieved the highest accuracy in CR sentiment task, and on par with skip-th and NBSVM in TREC, and it slightly underperformed the averaging models in Newsgroup categorization. In the next section, we analyze the Newsgroup and TREC datasets to shed light on intrinsic characteristics that may explain some of the performance variance.

6 Qualitative Analysis

Figure 3 shows t-SNE visualizations of the Newsgroup datasets using the various compositional models, including random and binary vectors. While random and binary vectors could identify shallow similarities between sentences as in STS tasks, they failed to do so in a globally cohesive manner. The random vectors also introduced noise in the representations, which resulted in a rather uniform vector space. All other models, except skip-th, clearly separated at least three regions that correspond to the categories **sport**, **computer**, and **religion/politics**. Smaller clusters with consistent labeling can also be identified with minimal separation between the clusters.

Table 6 shows examples of nearest neighbors using some of the models. skip-th vectors seem to be clustered more by structure than semantic content, unlike the doc2vec and sif models. To quantify these differences, we applied k-means clustering using $k = 3$ and $k = 8$, and calculated the clustering purity for each model as follows:

$$P(C, L) = \frac{1}{N} \sum_k \max_j |c_k \cap \ell_j| \quad (4)$$

where $C = \{c_1, \dots, c_K\}$ is the set of clusters, $L = \ell_1, \dots, \ell_K$ is the set of labels, and N the total number of samples. As shown in Table 7, using doc2vec and the averaging models, including GRAN, k-means successfully separated the 3 categories, with doc2vec and sif outperforming in both the fine-grained and coarse clustering. skip-th clusters, on the other hand, did not correspond with the correct labels, underperforming binary and random features, which explains its relatively low performance in text categorization tasks. inferSent achieved higher purity than binary, random and skip-th vectors but lower than the other models, particularly with $k = 3$.

sif †	And they work especially well when the Feds have cut off your utilities . The Dividians did not have that option after the FBI cut off their electricity . Not when the power has been cut off for weeks on end.	What does this bill do? And Bill James is not? Do you own " the Bill James players rating book"? Who has to consider it ? The being that does the action? I am still not sure I know what you are trying to say.
doc2vec	And they work especially well when the Feds have cut off your utilities . The Dividians did not have that option after the FBI cut off their electricity . Can the Feds get him on tax evasion ? I do not remember hearing about him running to the Post Office last night.	I did not claim that our system was objective . Did I claim that there was an absolute morality , or just an objective one? I have just spent two solid months arguing that no such thing as an objective moral system exists.
skip-th	What does this bill do ? Where do I get hold of these widgets ? What gives the United States the right to keep Washington D.C.? What makes you think Buck will still be in New York at year 's end with George back?	I have just spent two solid months arguing that no such thing as an objective moral system exists. The amount of energy being spent on one lousy syllogism says volumes for the true position of reason in this group. I just heard this week that he has started on compuserve flying models forum now.

Table 6: Examples of nearest neighbors in the 20-Newsgroup dataset. †sif using si-skip.

Model	Newsgroup		TREC
	k = 3, C=3	k = 8, C=8	k=6, C=6
Random	0.5563	0.2431	0.4424
Binary	0.6236	0.2963	0.444
avg †	0.8465	0.3969	0.4481
sif †	0.8776	0.4523	0.4037
doc2vec	0.8625	0.4967	0.3855
skip-th	0.4471	0.1854	0.3896
GRAN	0.8227	0.3553	0.3514
inferSent	0.6562	0.3801	0.4424

Table 7: Clustering purity measure with coarse categories (sports, computers, religion/politics) and the original 8 categories for the Newsgroup dataset, and 6 categories for TREC. †avg and sif using si-skip.

Figure 4 shows t-SNE visualizations of the questions in TREC training set. While all models identified some of the categories, like **HUM** and **LOC**, the `skip-th` and binary vectors appears to be more cohesively clustered by type than the other models. The question types are scattered in multiple smaller clusters, however, which explains why k-means clustering resulted in lower purity scores than `doc2vec` and averaging with $k = 6$. In Figure 5, purity results with various k are plotted. While purity is expected to increase with larger k , the rate of increase is much higher for `skip-th` than all other models, including binary features. This is consistent with the t-SNE visualization which shows several consistent clusters with `skip-th` that are larger than the binary clusters. `inferSent`'s performance was on par with `avg`, which is slightly lower than binary and `skip-th`, although the performance in the supervised setting was equivalent.

Table 8 shows nearest neighbors to the question "What country do the Galapagos Islands belong to?" using the various models. The averaging model clustered questions about islands; we observed similar behavior using weighted averaging, `doc2vec` and GRAN. On the other hand, `skip-th` clustered questions that start with "what country", which happens to be more suitable for identifying the **LOC** question type. Using binary vectors, questions that include the words "What" and "country" were clustered together, which do not necessarily correspond to the same question type. `inferSent` vectors seem to be clustered by a combination of semantic and syntactic features.

6.1 Discussion and Conclusions

In this study, we attempted to identify qualitative differences among the compositional models and general characteristics of their vector spaces that explain their performance in downstream tasks. Identifying the specific features that are most useful for each task may shed light on the type of information they encode and help optimize the representations for our needs. Word vector averaging performed reasonably well in most supervised benchmarks, and weighted averaging resulted in better performance in unsupervised similarity tasks outperforming all other models. Using the subword skipgram model for word embeddings resulted in better representations overall, particularly with `sif` weighting. The only model that performed on par with or slightly better than weighted averaging in unsupervised STS was

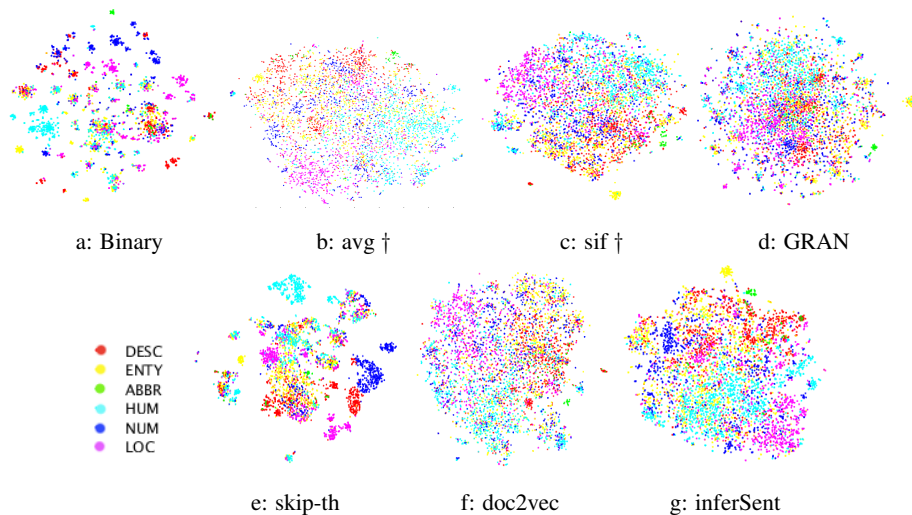


Figure 4: t-SNE visualizations of vectors on the TREC dataset. † avg and sif using si-skip.

avg†	What currents affect the area of the Shetland Islands and Orkney Islands in the North Sea?
	What two Caribbean countries share the island of Hispaniola?
Binary	What is a First World country?
	What is the best college in the country?
skip-th	What country is the worlds leading supplier of cannabis?
	What country did the Nile River originate in?
	What country boasts the most dams?
inferSent	What country did the ancient Romans refer to as Hibernia?
	How many islands does Fiji have?
	What country does Ileana Cotrubas come from ?

Table 8: Nearest neighbors to “What country do the Galapagos Islands belong to?” in TREC. †avg using si-skip

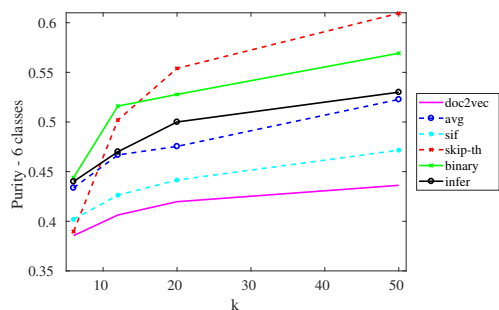


Figure 5: Clustering purity with increasing k on TREC. [avg and sif for si-skip]

inferSent. All models achieved higher correlation scores in the supervised STS evaluation, including skip-th, which performed poorly in the unsupervised setting. This suggests that at least some of the features in the vector space encode the semantic content and the remaining features are superfluous or encode structural information.

doc2vec and GRAN representations were qualitatively similar to idf and sif vectors where sentences/paragraphs were clustered by topic and semantic similarity. skip-th vectors, on the other hand, seemed to prominently represent structural rather than semantic features, which makes them more suitable for supervised tasks that rely on sentence structure rather than unsupervised similarity or topic categorization. inferSent vectors performed consistently well in all evaluation benchmarks, and a qualitative analysis of the vector space suggests that the vectors encode a balance of semantic and syntactic features. This makes inferSent suitable as a general-purpose model for sentence representation, particularly in supervised classification. For topic categorization, none of the compositional models outperformed the NBSVM baseline, which achieved significantly higher accuracies in all supervised topic categorization tasks. However, the distributional models, particularly weighted averaging, are more suitable in unsupervised or low-resource settings since sentences tend to be clustered cohesively by topic similarity and semantic relatedness.

References

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association of Computational Linguistics*.

- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval 2017)*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- Bill Dolan, Chris Quirk, and Chris Brockett. 2004. Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources. In *Proceedings of the 20th international conference on Computational Linguistics*.
- Yvette Graham and Timothy Baldwin. 2014. Testing for significance of increased correlation with human judgment. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 172–176.
- Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on World Wide Web*.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. Learning distributed representations of sentences from unlabelled data. pages 1367–1377.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*.
- Yangfeng Ji and Jacob Eisenstein. 2013. Discriminative improvements to distributional sentence similarity. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 891–896.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- Jey Han Lau and Timothy Baldwin. 2016. An empirical evaluation of doc2vec with practical insights into document embedding generation. *arXiv preprint arXiv:1607.05368*.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International Conference on Machine Learning*, pages 1188–1196.
- Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, Roberto Zamparelli, et al. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*, pages 216–223.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *ICLR*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Dmitrijs Milajevs, Dimitri Kartsaklis, Mehrnoosh Sadrzadeh, and Matthew Purver. 2014. Evaluating neural word representations in tensor-based compositional settings. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd annual meeting on Association for Computational Linguistics*, page 271.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 115–124.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Benjamin Riedel, Isabelle Augenstein, Georgios P Spithourakis, and Sebastian Riedel. 2017. A simple but tough-to-beat baseline for the fake news challenge stance detection task. *arXiv preprint arXiv:1707.03264*.
- Tobias Schnabel, Igor Labutov, David Mimno, and Thorsten Joachims. 2015. Evaluation methods for unsupervised word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 298–307.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Sida Wang and Christopher D Manning. 2012. Baselines and bigrams: Simple, good sentiment and topic classification. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*, 39(2-3):165–210.
- John Wieting and Kevin Gimpel. 2017. Revisiting recurrent networks for paraphrastic sentence embeddings.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015a. Towards universal paraphrastic sentence embeddings. *ICLR*.
- John Wieting, Mohit Bansal, Kevin Gimpel, Karen Livescu, and Dan Roth. 2015b. From paraphrase database to compositional paraphrase model and back. *Transactions of the Association for Computational Linguistics*.