

Original Paper

Word Embedding Dimension Reduction via Weakly-Supervised Feature Selection

Jintang Xue*, Yun-Cheng Wang, Chengwei Wei and C.-C. Jay Kuo

University of Southern California, Los Angeles, California, USA

ABSTRACT

As a fundamental task in natural language processing, word embedding converts each word into a representation in a vector space. A challenge with word embedding is that as the vocabulary grows, the vector space's dimension increases, which can lead to a vast model size. Storing and processing word vectors are resource-demanding, especially for mobile edge-devices applications. This paper explores word embedding dimension reduction. To balance computational costs and performance, we propose an efficient and effective weakly-supervised feature selection method named WordFS. It has two variants, each utilizing novel criteria for feature selection. Experiments on various tasks (e.g., word and sentence similarity and binary and multi-class classification) indicate that the proposed WordFS model outperforms other dimension reduction methods at lower computational costs.

Keywords: Word embedding, feature selection, dimension reduction, natural language processing

*Corresponding author: Jintang Xue, jintangx@usc.edu

1 Introduction

In recent years, large language models (LLMs) have made a breakthrough in natural language processing (NLP). These models have revolutionized the understanding and generation of human language tasks such as question-answering, machine translation, and text summarization [30, 31, 17], etc. As a fundamental component in language models, word embedding represents words as vectors in a continuous, high-dimensional space [29]. Word vectors capture semantic and syntactic meanings of words, providing word relationships for various downstream tasks. Static word embedding methods, such as Glove [18], Word2Vec [14], and Fasttext [2], assign a fixed vector to each word. They convert input words to corresponding vectors for further processing and model training. Contextual word embedding methods, such as ELMo [19] and BERT [4], leverage deep-learning models to generate word vectors based on the context. A word can have a different contextual word embedding in a different sentence, allowing the model to capture subtle differences in the meaning of the same word in various contexts.

A common challenge for static and contextual word embedding methods is that the high dimension of a word vector leads to an enormous model size. Typically, a word vector has hundreds to thousands of dimensions. For instance, storing a vocabulary of 3 million words in 300 dimensions would require 3.39 GB. Loading a 300-dimensional word embedding matrix with 2.5 million tokens would require up to 6 GB of memory on a 64-bit system [20]. On the one hand, high-dimensional word vectors provide a good representation of complex human language, which is crucial for performing downstream tasks. On the other hand, high-dimensional word vectors have higher demands on computational resources and memory requirements. Thus, dimension reduction is critical in the application of word vectors.

Existing dimension reduction methods mainly consist of traditional PCA-based and deep-learning-based models. For PCA-based models, [20] combines a post-processing technique with PCA to achieve an effective method for dimension reduction. For deep-learning-based models, [8] proposed a deep-learning method called EmbedTextNet for word embedding dimension reduction by leveraging a VAE model with a correlation penalty added to the weighted reconstruction loss. The model works well in low-dimensional embedding sizes but takes a long time to train. Other deep learning methods focus on model compression [10, 22] and quantization [11, 24] rather than dimension reduction.

Traditional unsupervised PCA-based methods are known for their efficiency and interpretability. In contrast, supervised deep-learning-based methods are inefficient and lack interpretability. For instance, autoencoders take longer training and inference time in sentence embedding dimension reduction than traditional methods [33]. Semi-supervised feature selection methods can also be used for dimension reduction. Since word vectors can be viewed as extracted

features for each word, feature selection methods have the potential to provide a straightforward yet effective way to reduce the word dimension. It could help balance unsupervised PCA-based methods and resource-intensive deep-learning-based methods.

This paper investigates using a small subset of labeled word similarity pairs to develop a weakly supervised dimension reduction method called WordFS (i.e., word dimension reduction with **F**eature **S**election). We demonstrate that one can achieve dimension reduction effectively by supervising a limited number of word similarity pairs. Note that word similarity is typically used in evaluation tasks but is rare in word embedding dimension reduction, a key novelty of this work. The proposed WordFS method consists of three stages: 1) post-processing, 2) feature extraction, and 3) weakly-supervised feature selection. We apply WordFS to other downstream tasks to show its generalizability. Experimental results show that WordFS outperforms existing methods in word similarity and various tasks while achieving much lower computational costs.

This work has the following significant contributions.

- We propose a novel, effective, and efficient dimension reduction method called WordFS for word embeddings from the perspective of feature selection based on weakly-supervised learning.
- We demonstrate the potential of combining feature selection methods and word similarity for word embedding dimension reduction.
- We show the effectiveness and efficiency of our approach on various downstream tasks, including sentence similarity and classification tasks. Our method generally outperforms the existing techniques while being more straightforward and efficient.

The rest of the paper is organized as follows. Related work is reviewed in Section 2. The proposed WordFS method is described in Section 3. Experimental results are shown in Section 4. Finally, concluding remarks and future extensions are given in Section 5.

2 Related Work

Word embedding compression is an essential topic for storing and processing word embeddings, especially on computationally limited devices. Existing dimension reduction methods mainly consist of traditional and deep learning-based models.

2.1 Matrix Decomposition Techniques

Matrix decomposition techniques, such as singular value decomposition (SVD), principal components analysis (PCA), non-negative matrix factorization (NMF), and factor analysis (FA), have been applied to dimension reduction of word embeddings. While simple and efficient, these methods are not highly effective.

Post-processing methods help enhance word embeddings in dimension reduction. Effective construction of lower dimensional word embeddings can be achieved by combining post-processing algorithms (PPAs) based on reducing the anisotropy [16] with the PCA method [20]. The dimension-reduced vectors can perform similarly or even better than the original pre-trained word embeddings, outperforming most unsupervised methods. [9] extends this method to contextual embeddings by adding in additional geodesic distance information via the Isomap algorithm [25]. PCA also performs well as an unsupervised dimension reduction method for pre-trained sentence embeddings [33].

2.2 Non-Linear Dimensionality Reduction Techniques

Non-linear dimensionality reduction techniques like uniform manifold approximation and projection (UMAP) [13] and t-distributed stochastic neighbor embedding (t-SNE) [27] have been used on dimension reduction tasks. UMAP is a dimensionality reduction method based on manifold theory and topological data analysis. It constructs a topological representation of high-dimensional data by approximating local manifolds and stitching together their fuzzy simplicial sets. However, UMAP’s performance in reducing the dimensionality of word embeddings is not satisfactory. t-SNE is a statistical method based on stochastic neighbor embedding. It assigns each data point a location in a two- or three-dimensional map for visualizing high-dimensional data. Applying t-SNE to a large corpus while keeping many dimensions requires a significant amount of time and memory, making it unsuitable for this task.

2.3 Deep Learning Methods

Deep learning-based models have recently become quite popular, and efforts have been made to explore their potential in reducing word embedding dimensions. EmbedTextNet [8] achieves better performance in low-dimensional embedding sizes by utilizing a VAE model with a correlation penalty added to the weighted reconstruction loss. However, deep learning models typically require more computational resources and longer training time, contradicting the original goal of dimension reduction.

2.4 Feature Selection Methods

Feature selection is another simple and intuitive method for removing redundant features. Feature selection methods can be applied to bag-of-words representations, where each unique word is treated as a distinct feature, for emotion recognition [5] and sentiment classification [26] tasks to reduce the number of terms and save storage and memory space. Also, feature selection methods can be applied for text classification to select the most representative word [21]. However, the role of feature selection for word embedding dimension reduction is under-explored. This paper shows that a straightforward and intuitive feature selection method can effectively reduce dimension. We believe applying the feature selection method for word embedding dimension reduction is novel.

3 Proposed Method

3.1 System Overview

Our WordFS method consists of three stages: 1) post-processing, 2) feature extraction, and 3) feature selection, as shown in Figure 1. We adopt word similarity datasets as the source of supervision to guide the feature selection in our model. As a widely-used benchmark of word embeddings, the word similarity task evaluates the quality of word embeddings in representing semantic and syntactic meanings by measuring how closely the representation of word vectors matches human perception of similarity [28], making them good supervision resources for word embedding dimension reduction. Also, since word similarity captures fundamental characteristics of word embeddings, which are crucial for many NLP tasks, its potential for generalizing to downstream tasks is noteworthy.

In our method, We begin by optionally applying a post-processing method to our pre-trained word embeddings since the anisotropy may not always be harmful [1]. Word vectors pre-trained on a smaller amount of data may be susceptible to noise and could benefit from post-processing. On the other hand, word vectors pre-trained on a larger corpus of data may capture important nuanced information necessary for downstream tasks but could be negatively impacted by post-processing. We then extract pair-wise features for each word pair by conducting element-wise production with normalization. Then, we adopt two different criteria for feature selection to evaluate each dimension and identify essential dimensions. The first one is based on a supervised feature selection method called RFT [32]. We select mean squared error (MSE) as our cost function. The second is based on Spearman’s rank correlation coefficient between each feature dimension and word similarity scores. Finally, the selected dimensions of the word embeddings are kept, and other dimensions

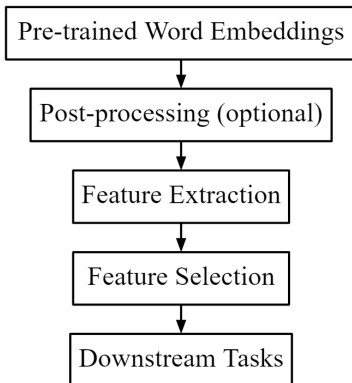


Figure 1: An overview of the proposed WordFS method.

are discarded according to the requirements for dimension reduction. The details of each stage are elaborated below.

3.2 *Post-processing*

To improve the quality of word embedding in downstream applications, performing post-processing techniques on pre-trained word embeddings can be crucial. The post-processing algorithm (PPA) [16] is an effective post-processing method to improve the isotropy of word representations by removing the common mean vector and the top principal components of all words. Details of PPA are shown in Algorithm 1 [16]. First, the mean vector is computed from all embeddings in the vocabulary and then subtracted from each embedding. Next, the PCA components are derived from the adjusted mean matrix, with the top-D components being chosen. Finally, the projections on the top-D principal components are removed from each word embedding.

The existing effective PCA-based method [20] applies the PPA before and after dimension reduction to enhance both the original and the dimension-reduced word vectors. In our method, we only apply PPA before dimension reduction, which makes our method simpler. Applying PCA-based post-processing to the selected subset after feature selection may disrupt the previous results because PCA and the feature selection criteria can have different objectives. Specifically, PCA transforms features to maximize variance, while feature selection methods usually select features based on their correlation with labels. Also, the selected feature subset may lose crucial information for PCA to find valuable principal components. As an enhancement of original word embeddings before dimension reduction, we also make PPA an optional choice. The reason is that anisotropy may not always be harmful [1], and using

Algorithm 1 Post-processing Algorithm (PPA)

- 1: **Input:** Word vector matrix M ; Threshold D
 - 2: Subtract the mean:
 $n \leftarrow$ number of word vectors in M
 $\mu \leftarrow \frac{1}{n} \sum_{i=1}^n M_i$
for each word vector M_i **in** M **do**
 $M_i \leftarrow M_i - \mu$
end
 - 3: Compute the PCA components:
 $U \leftarrow PCA(M)$
for $i = 1$ **to** D **do**
 $u_i \leftarrow$ the i^{th} principal component from U
end
 - 4: Remove the top- D principal components:
for all v in M **do**
 $v' \leftarrow v - \sum_{i=1}^D (u_i^\top v) u_i$
end
 - 5: **Output:** Enhanced word vector matrix M'
-

PPA depends on the pre-trained word embeddings and specific application scenarios. Specifically, word embeddings trained on less data may be vulnerable to noise. Applying PPA helps mitigate some of the noise and offers better performance. On the other hand, using PPA may result in the loss of crucial and intricate information related to word representation, which can harm performance. Word embeddings trained on relatively small data benefit from uniformly distributed word embeddings, while complex contextual tasks may require anisotropy to capture detailed information.

3.3 Feature Extraction

In this stage, we extract suitable features for feature selection based on the dataset that provides supervision. In this paper, we leverage word similarity data as the weak supervision of our model. Since the word similarity is based on the correlation (i.e., word similarity score) between every two words, we extract pair-wise features for each word similarity pair.

The pair-wise features are constructed in each dimension based on cosine similarity, an effective measure for evaluating the similarity between word embeddings. Cosine similarity between two-word vectors is derived by dividing their dot product by the product of their magnitudes:

$$\begin{aligned} \text{Sim}(\mathbf{a}, \mathbf{b}) &= \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} \\ &= \frac{\sum_{i=1}^n a_i b_i}{\sqrt{\sum_{i=1}^n a_i^2} \cdot \sqrt{\sum_{i=1}^n b_i^2}}, \end{aligned} \quad (1)$$

where \mathbf{a} and \mathbf{b} are embeddings of the two words, and n represents the number of elements in a word embedding. a_i and b_i are the i th elements of \mathbf{a} and \mathbf{b} , respectively. The cosine similarity score can be viewed as the sum of the normalized results of element-wise products. Each pair of elements (i.e., each dimension) has a linear impact on the cosine similarity score of that pair of word vectors. Therefore, the normalized results of element-wise products can be viewed as features to predict the similarity score. The feature we use for each dimension is:

$$f_i = \frac{a_i b_i}{\sqrt{\sum_{i=1}^n a_i^2} \cdot \sqrt{\sum_{i=1}^n b_i^2}}, \quad (2)$$

where f_i is the feature of the i^{th} dimension extracted from the corresponding elements a_i and b_i of the two word embeddings.

3.4 Feature Selection

The method for the feature selection process depends on the objective of downstream tasks. Specifically, we employ two feature selection methods for prediction and similarity tasks. These two feature selection methods are developed into our model’s variants, WordFS-S and WordFS-P. We focus on filter feature selection techniques since they are simple and computationally efficient [23]. Although different dimensions of distributed word embeddings may capture nuanced and interdependent semantic relationships, they allow us to identify dimensions that contribute most significantly to model performance. Besides, filter feature selection techniques select features based on their correlation with the labels and are independent of the classifier. As a result, the selected features are generalizable.

Prediction tasks: For prediction tasks, we leverage the methods from Discriminant Feature Test (DFT) and Relevant Feature Test (RFT) [32] to build our WordFS-P model. DFT and RFT are a pair of filter feature selection techniques proposed recently for classification and regression tasks, respectively. They have been widely used in green learning architectures [12] for feature dimension reduction to reach a smaller model size. In our experiments, RFT is utilized for feature selection, and we adopt the word similarity scores as the labels. RFT partitions each feature dimension into two subintervals and calculates the overall mean square error (MSE) as the loss function. A smaller loss function means a better feature dimension. Specifically, given a feature of the i^{th} dimension f_i , the feature space is partitioned into $B = 2^k$, $k = 1, 2, \dots$

uniform segments and the optimal partition threshold is searched among $B - 1$ candidates in the range $[\min(f_i), \max(f_i)]$:

$$f_b^i = \min(f_i) + \frac{b}{B} [\max(f_i) - \min(f_i)], \quad (3)$$

where $b = 1, 2, \dots, B - 1$. A threshold t partitions the i^{th} feature space into the left subset $S_{L,t}^i$ and the right subset $S_{R,t}^i$. The MSEs for regression $R_{L,t}^i$ and $R_{R,t}^i$ are separately calculated in the two subsets by using the mean of target values as the estimated regression value of all samples. Then, the RFT loss is defined as the weighted sum of the two MSEs:

$$R_t^i = \frac{N_{L,t}^i R_{L,t}^i + N_{R,t}^i R_{R,t}^i}{N}, \quad (4)$$

where $N_{L,t}^i$ and $N_{R,t}^i$ denote the number of samples in each subset and N is the total number of samples. The RFT loss of the i^{th} feature dimension is defined as the minimal RFT loss among all the candidate thresholds:

$$R_i = \min_{t \in T} R_t^i. \quad (5)$$

Finally, all feature dimensions' RFT loss is ranked in ascending order, and the top K dimensions are selected, where K is the dimension of the word vectors after our dimension reduction approach.

Similarity tasks: Similarity tasks are usually done by calculating the cosine similarity between target vectors, which differs from prediction tasks. The cost function of DFT/RFT may not match well with the evaluation criteria of similarity tasks. Inspired by the evaluation of word similarity, we use Spearman's rank correlation coefficient to develop our WordFS-S model. It involves finding the correlation between the ranks generated by features extracted from each dimension and the target labels. First, the ranked values of each feature dimension and the labels are calculated. Second, Spearman's rank correlation coefficient for each feature dimension is calculated by the Pearson correlation coefficient between the ranked values of each dimension and the labels:

$$r_i = \rho_{R(f_i), R(y)} = \frac{\text{cov}(R(f_i), R(y))}{\sigma_{R(f_i)} \sigma_{R(y)}}, \quad (6)$$

where $R(\cdot)$ represents the rank of given raw scores. The higher the correlation coefficient, the better the feature dimension is. Finally, the Spearman's rank correlation coefficients of all the feature dimensions are then ranked in descending order, and the top K dimensions are selected.

4 Experiments

We evaluate our method by applying it to several pre-trained word embeddings. The word similarity datasets provide weak supervision to guide our feature selection module. The dimension-reduced word embeddings are used for word similarity tasks and downstream tasks. Then, we compare the results of our method with the original pre-trained word embeddings, the reduced word embeddings from UMAP [13], the dimension-reduced word embeddings from the PCA-based method called Algo [20], and the results from the deep learning-based method called EmbedTextNet [8]. The performance of different datasets can vary due to their specific requirements and emphasis. Finding a general dimensionality reduction method that performs well across all datasets without fine-tuning is challenging. Therefore, we focus on comparing the average performance for each task. The method that achieves better performance on average is considered better.

4.1 Pre-trained Word Embeddings

In our experiments, we use the following three pre-trained word embeddings.

1. Glove word embeddings [18] trained on Wikipedia 2014 and Gigaword 5 corpus (6B tokens, 400K vocabulary)
They are available in multiple dimensions, precisely 50, 100, 200, and 300.
2. Word2vec word embeddings [14] trained on a portion of the Google News dataset (about 100B words)
The model provides word vectors of 300 dimensions for 3M words and phrases.
3. Fasttext word embeddings [2] trained on Wikipedia 2017, UMBC webbase corpus, and statmt.org news dataset (16B tokens) [15]
It generates 1 million 300-dimensional word vectors. The Gloved-based model is trained on the least amount of data among the three pre-trained word embeddings.

4.2 Word Similarity Datasets

We use the widely-used word similarity datasets [6] to evaluate our method. The word similarity datasets contain word pairs and their corresponding scores from human annotators based on perceived relatedness or similarity. The cosine similarity of each pair of words calculates the similarity score from the word embeddings. We use Spearman’s rank correlation coefficient as our evaluation metric. It measures how closely the ranking derived from the cosine

similarities of given word vectors matches those based on human judgments. A higher value of this metric indicates a better match to human-labeled similarity rankings.

We first evaluate our methods on twelve word similarity datasets based on all the three word embeddings. In later experiments, we will assess the generalizability of our WordFS model across various embeddings on multiple downstream tasks. Since we choose word similarity as the guidance for feature selection, we apply 5-fold cross-validation to each word similarity dataset to train and test our method. The reported results are the average Spearman’s rank correlation coefficients from the five folds. We evaluate the pre-trained word embeddings and all the methods on the same cross-validation fold and take the average for a fair comparison. All the results we report average over five different cross-validation trials in all the experiments. We set the number of bins in RFT to 4, the default and recommended value from the RFT model, and the threshold in the PPA to 7, the same as in the PCA-based model.

Table 1 shows the results across twelve word similarity datasets of different dimension reduction methods reduced from 300-dimensional Glove word embeddings to 150, 100, and 50 dimensions, respectively. In the table, WordFS-P and WordFS-S represent our proposed methods based on RFT and similarity feature selection methods, respectively. P represents PPA, and wP and woP represent PPA and without PPA. Algo is the method proposed in [20], which utilizes PPA before and after PCA. The last column is the average of the previous columns. The UMAP results are significantly lower compared to other methods. This is understandable because UMAP is a general dimension reduction method and is not specifically designed to reduce the dimension of word embeddings, unlike other methods. For the dimension-reduced vectors with dimensions of 50, 100, and 150 based on the Glove embedding, our methods outperform the original Glove word embeddings and other methods in most datasets across these reduced dimensions. Our WordFS-S-wP method achieves the highest average score when the dimension is reduced to 50 and 150 and is close to the highest, which is our WordFS-P-wP method, with a difference of 0.03 when reduced to 100. Both feature selection-based approaches perform better than the existing methods. As expected, the WordFS-S-wP method performs better than the WordFS-P-wP method in this task. Our WordFS-S-wP method outperforms the best performance from existing methods with an average improvement of 0.61, 1.97, and 3.69 in terms of Spearman’s rank correlation coefficients when reducing to 150, 100, and 50 dimensions, respectively. Our method can perform better than the original 300-dimensional word embeddings, even when reduced to 100 dimensions. For the dimension-reduced vectors based on the Word2vec embeddings, applying PPA may decrease performance, as the word embeddings are pre-trained on larger data, containing subtle linguistic connections. Our method can perform highest when the word vector dimension is reduced to 150 and 100 and slightly lower when reduced

Table 1: Comparison of Spearman’s rank correlation coefficient across multiple word similarity datasets, where **Boldface** indicates the best value in each column, and underline indicates the second best value in each column.

| Dataset | MC -30 | Men -TR -3k | M Turk -287 | M Turk -771 | RG -65 | RW Stan ford | Sim Lex -999 | VE RB -143 | WS -353 -ALL | WS -353 -REL | WS -353 -SIM | YP -130 | Avg |
|-------------------------|--------------|-------------------|-------------------|-------------------|--------------|--------------------|--------------------|------------------|--------------------|--------------------|--------------------|--------------|--------------|
| Glove-300D | 64.03 | 73.71 | 62.44 | 64.60 | 72.97 | 41.16 | 36.97 | 29.97 | 60.27 | 55.44 | 64.74 | 55.15 | 56.79 |
| UMAP-150D | 25.39 | 38.81 | 35.88 | 31.54 | 32.50 | 13.91 | 25.33 | 24.13 | 30.76 | 20.19 | 38.23 | 14.00 | 27.56 |
| Algo-150D | 67.97 | 75.21 | 62.47 | 64.11 | <u>72.94</u> | 43.45 | 37.81 | 38.06 | 66.71 | <u>60.55</u> | <u>70.74</u> | 58.09 | 59.84 |
| EmbedTextNet-150D | 79.59 | 75.43 | 58.36 | 62.91 | 80.00 | 41.31 | 37.49 | 27.33 | 63.06 | 55.35 | 67.50 | 55.80 | 58.68 |
| WordFS-P-woP-150D(ours) | 66.79 | 74.94 | 60.68 | 63.44 | 69.83 | 36.16 | 37.93 | 30.26 | 62.49 | 53.69 | 66.78 | 52.27 | 56.27 |
| WordFS-P-wP-150D(ours) | <u>73.63</u> | <u>75.38</u> | <u>64.38</u> | <u>64.24</u> | 70.37 | <u>45.01</u> | 42.47 | 34.88 | 68.53 | 60.58 | 70.25 | 50.54 | <u>60.02</u> |
| WordFS-S-woP-150D(ours) | 67.79 | <u>75.78</u> | 62.62 | 62.95 | 70.26 | 36.07 | <u>44.72</u> | <u>36.61</u> | 64.72 | 55.36 | 69.42 | 45.76 | 57.67 |
| WordFS-S-wP-150D(ours) | 72.27 | 75.96 | 65.57 | 64.35 | 70.69 | 45.22 | 45.08 | 36.03 | <u>67.16</u> | 59.18 | 71.00 | <u>52.90</u> | 60.45 |
| Glove-100D | 57.18 | 68.08 | 60.78 | 57.81 | 65.29 | 36.54 | 29.79 | 30.61 | 52.83 | 47.62 | 58.55 | 43.02 | 50.68 |
| UMAP-100D | 22.88 | 39.72 | 38.15 | 31.06 | 40.10 | 14.97 | 26.08 | 18.97 | 30.78 | 21.48 | 37.27 | 14.98 | 28.04 |
| Algo-100D | 69.99 | 70.64 | 60.65 | 60.92 | 75.93 | 39.55 | 35.80 | 29.70 | 62.33 | 53.29 | 68.49 | 50.59 | 56.49 |
| EmbedTextNet-100D | <u>70.21</u> | 70.43 | <u>62.04</u> | 59.28 | <u>74.42</u> | 37.20 | 35.19 | 30.63 | 59.66 | 54.10 | 64.94 | 44.79 | 55.24 |
| WordFS-P-wP-100D(ours) | 72.91 | <u>73.96</u> | 61.16 | <u>60.54</u> | 69.91 | 43.65 | <u>42.65</u> | <u>34.93</u> | 67.80 | 58.91 | <u>69.15</u> | <u>46.28</u> | 58.49 |
| WordFS-S-wP-100D(ours) | 69.63 | 74.72 | 65.29 | 60.23 | 68.79 | <u>43.16</u> | 43.95 | 37.17 | <u>66.11</u> | <u>56.79</u> | 69.65 | 46.06 | <u>58.46</u> |
| Glove-50D | 53.75 | 65.24 | 60.67 | 55.14 | 58.41 | 34.02 | 26.44 | 25.26 | 49.82 | 44.94 | 55.66 | 36.03 | 47.11 |
| UMAP-50D | 17.39 | 39.84 | 36.15 | 31.14 | 41.56 | 15.13 | 26.22 | 22.71 | 31.37 | 20.87 | 38.87 | 8.59 | 27.49 |
| Algo-50D | <u>64.48</u> | 61.43 | 53.19 | 47.30 | <u>63.65</u> | 33.65 | 26.75 | 34.77 | 55.60 | 45.01 | 59.35 | 38.00 | 48.60 |
| EmbedTextNet-50D | 55.58 | 65.90 | 63.61 | 55.18 | 53.97 | 32.60 | 28.57 | <u>32.21</u> | 57.51 | 50.19 | 64.01 | 36.18 | 49.63 |
| WordFS-P-wP-50D(ours) | 63.11 | <u>69.91</u> | 55.14 | <u>52.58</u> | 63.52 | 39.00 | 41.33 | 30.49 | 64.35 | <u>53.76</u> | <u>64.28</u> | <u>37.64</u> | <u>52.92</u> |
| WordFS-S-wP-50D(ours) | 67.81 | 71.06 | <u>58.97</u> | 52.08 | 64.96 | <u>36.62</u> | <u>39.96</u> | 30.28 | <u>60.46</u> | <u>56.47</u> | 66.33 | 34.86 | 53.32 |
| Word2vec-300D | 73.42 | 77.01 | 68.52 | 66.73 | 71.65 | 53.39 | 44.19 | 49.69 | 68.92 | 61.42 | 76.87 | 53.80 | 63.80 |
| UMAP-150D | 36.38 | 50.43 | 42.53 | 26.17 | 35.13 | 26.00 | 22.32 | 6.25 | 42.09 | 32.59 | 51.84 | 17.37 | 32.42 |
| Algo-150D | 78.91 | 78.23 | <u>63.46</u> | 65.86 | 74.93 | 51.43 | 42.53 | 41.13 | 68.25 | <u>60.28</u> | 75.11 | 48.34 | 62.37 |
| EmbedTextNet-150D | 72.28 | 76.33 | 63.09 | <u>65.28</u> | <u>72.33</u> | <u>52.27</u> | 43.53 | 44.47 | 66.44 | 57.95 | <u>75.04</u> | 48.89 | 61.49 |
| WordFS-P-woP-150D(ours) | 68.16 | 76.16 | 63.35 | 63.39 | 70.79 | 51.60 | 45.32 | 59.99 | 67.68 | 58.04 | 72.44 | 48.14 | 61.75 |
| WordFS-P-wP-150D(ours) | 66.80 | <u>76.77</u> | 63.43 | 63.45 | 69.99 | 52.00 | 46.57 | 50.77 | 66.81 | 56.57 | 72.64 | 48.75 | 61.21 |
| WordFS-S-woP-150D(ours) | <u>75.71</u> | 75.93 | 64.22 | 63.35 | 71.63 | 51.86 | 46.95 | <u>51.89</u> | 69.58 | 61.30 | 72.85 | 53.07 | 63.19 |
| WordFS-S-wP-150D(ours) | 74.80 | 76.73 | 61.72 | 63.74 | 70.55 | 52.59 | <u>46.77</u> | 49.56 | <u>68.79</u> | <u>58.40</u> | 74.73 | <u>51.99</u> | <u>62.53</u> |
| UMAP-100D | 33.18 | 51.38 | 42.27 | 27.97 | 33.24 | 26.86 | 23.03 | 8.13 | 42.56 | 31.03 | 53.53 | 11.16 | 32.03 |
| Algo-100D | 80.06 | 76.08 | 63.86 | 63.56 | 72.26 | 48.21 | 39.76 | 30.78 | <u>67.49</u> | 59.81 | 73.40 | 39.65 | 59.58 |
| EmbedTextNet-100D | <u>74.11</u> | <u>75.73</u> | <u>62.21</u> | <u>63.14</u> | <u>71.10</u> | 50.33 | 41.89 | 40.12 | 66.10 | <u>59.38</u> | <u>72.89</u> | <u>45.34</u> | <u>60.19</u> |
| WordFS-P-woP-100D(ours) | 62.45 | 74.14 | 57.35 | 61.32 | 69.55 | <u>50.01</u> | <u>45.18</u> | <u>52.28</u> | 64.82 | 55.96 | 71.10 | 44.29 | 59.04 |
| WordFS-S-woP-100D(ours) | 65.42 | 74.45 | 60.81 | 60.28 | 66.60 | 49.75 | 45.70 | 52.60 | 68.49 | 57.76 | 70.23 | 50.48 | 60.21 |
| UMAP-50D | 28.38 | 50.06 | 42.88 | 27.44 | 35.31 | 25.90 | 22.21 | 9.78 | 44.41 | 34.64 | 54.23 | 16.77 | 32.67 |
| Algo-50D | 73.87 | 70.36 | <u>62.29</u> | 54.36 | 66.97 | 41.89 | 32.36 | 22.72 | 59.75 | 53.79 | <u>66.67</u> | <u>37.28</u> | <u>53.53</u> |
| EmbedTextNet-50D | <u>72.49</u> | 73.10 | 66.33 | 59.24 | 62.16 | 46.30 | 36.71 | 26.65 | 63.99 | 55.74 | 72.62 | 35.26 | 55.88 |
| WordFS-P-wP-50D(ours) | 49.65 | 69.72 | 53.24 | <u>55.01</u> | <u>63.96</u> | 44.47 | <u>43.12</u> | <u>41.79</u> | 57.92 | 47.89 | 65.01 | 34.51 | 52.19 |
| WordFS-S-woP-50D(ours) | 54.68 | <u>70.87</u> | 51.66 | 54.62 | 55.86 | <u>44.65</u> | 43.52 | 47.26 | <u>61.96</u> | <u>49.05</u> | 65.16 | 38.93 | 53.18 |
| Fasttext-300D | 77.08 | 79.01 | 69.84 | 70.57 | 82.68 | 52.23 | 45.11 | 43.86 | 70.63 | 64.65 | 79.66 | 48.92 | 65.35 |
| UMAP-150D | 56.94 | 59.32 | 46.98 | 43.50 | 49.67 | 33.78 | 29.16 | 32.07 | 40.20 | 29.54 | 54.10 | 20.65 | 41.33 |
| Algo-150D | 92.63 | <u>80.59</u> | 71.32 | 70.74 | 87.97 | 51.11 | 45.04 | <u>45.03</u> | 73.18 | <u>68.03</u> | 77.21 | 49.31 | 67.68 |
| EmbedTextNet-150D | <u>77.54</u> | 78.50 | <u>68.39</u> | 69.78 | <u>82.10</u> | 52.03 | 45.71 | 38.11 | 68.01 | 62.98 | 74.77 | 51.57 | 64.12 |
| WordFS-P-woP-150D(ours) | 70.45 | 79.52 | 65.50 | 67.05 | 80.15 | 51.29 | 46.59 | 39.11 | 69.56 | 60.91 | 77.27 | 51.88 | 63.27 |
| WordFS-P-wP-150D(ours) | 76.17 | 80.13 | 67.26 | <u>70.29</u> | 79.83 | 54.58 | <u>47.44</u> | 38.84 | 72.93 | 66.76 | 78.53 | 55.68 | 65.70 |
| WordFS-S-woP-150D(ours) | 69.54 | 80.39 | 64.20 | 69.73 | 75.89 | 53.47 | 47.43 | 44.62 | <u>74.32</u> | 66.94 | 77.07 | 52.69 | 64.69 |
| WordFS-S-wP-150D(ours) | 73.42 | 80.89 | 66.50 | 68.87 | 78.78 | <u>53.94</u> | <u>49.14</u> | <u>46.14</u> | 74.99 | 71.09 | <u>77.40</u> | <u>53.71</u> | <u>66.24</u> |
| UMAP-100D | 55.58 | 59.31 | 46.16 | 43.69 | 41.73 | 33.21 | 28.80 | 32.09 | 39.17 | 29.49 | 53.00 | 19.40 | 41.33 |
| Algo-100D | 89.65 | 78.80 | 71.86 | <u>66.88</u> | 87.66 | 49.19 | 42.30 | 39.96 | 70.33 | 64.11 | 74.89 | 45.88 | 65.13 |
| EmbedTextNet-100D | <u>81.65</u> | 77.74 | <u>68.26</u> | 66.93 | <u>81.72</u> | <u>51.34</u> | 40.62 | <u>41.30</u> | 69.21 | 62.20 | <u>75.68</u> | 46.99 | 63.64 |
| WordFS-P-wP-100D(ours) | 70.91 | 77.40 | 63.64 | 66.07 | 76.81 | 52.08 | <u>46.48</u> | 35.23 | <u>71.05</u> | <u>65.77</u> | <u>74.08</u> | <u>52.72</u> | 62.69 |
| WordFS-S-wP-100D(ours) | 75.71 | <u>78.33</u> | 65.38 | 66.44 | 76.89 | 51.25 | 47.24 | 43.04 | 72.63 | 67.27 | 75.96 | 56.67 | <u>64.73</u> |
| UMAP-50D | 55.12 | 58.52 | 46.70 | 43.67 | 40.37 | 33.63 | 29.38 | 31.36 | 40.87 | 31.62 | 54.07 | 21.06 | 40.53 |
| Algo-50D | 79.37 | 72.74 | 70.66 | 61.06 | 77.50 | 45.47 | 35.79 | 37.48 | 59.66 | 52.72 | 67.79 | 43.79 | 58.67 |
| EmbedTextNet-50D | <u>74.80</u> | <u>72.39</u> | <u>69.21</u> | <u>60.69</u> | <u>73.40</u> | 45.50 | 36.04 | 42.57 | 70.13 | 62.00 | <u>76.17</u> | 34.52 | 59.78 |
| WordFS-P-wP-50D(ours) | 61.31 | 70.80 | 54.88 | 58.09 | 67.07 | 46.60 | <u>44.66</u> | 26.71 | 64.36 | <u>57.75</u> | 66.41 | <u>50.27</u> | 55.74 |
| WordFS-S-wP-50D(ours) | 67.94 | 71.97 | 62.42 | 60.33 | 67.92 | <u>45.53</u> | 45.48 | 41.89 | <u>67.41</u> | 59.24 | 70.99 | 54.36 | 59.62 |

to 50. For the Fasttext word embeddings, although our method may not be the best, it consistently performs very close to the highest performance.

Then, we aggregated the twelve word similarity datasets. We first scaled them into the same range $[0, 1]$ and then averaged the similarity scores of overlapped word pairs. There are two reasons to aggregate the datasets: 1) to provide an overall evaluation for the word similarity tasks and 2) to construct a comprehensive dataset to train our feature selection methods for various downstream tasks. We refer to our model as a weakly-supervised method for downstream tasks because there are only 7,705 human-labeled similarity scores for different word pairs in the aggregated dataset, which is much less than the total number of possible word pairs in the corpus to pre-train the word embeddings. For example, there are 400K word vectors in the pre-trained Glove word embeddings, which can generate approximately 80 billion unique word pairs. The other two word embeddings we use in the experiments have an even more extensive vocabulary. Compared with the possible word pairs, the guidance provided by the aggregated dataset is limited. We experimented with the aggregated dataset using the same settings as before.

Table 2 shows the results of various methods for dimension reduction on the aggregated dataset. The annotations in the tables are the same as in the previous one. The performance of UMAP is still very low. Our WordFS-S methods perform best among all the experiments with different pre-trained word embeddings in other dimensions. And there is a large gap between our methods and the existing methods. Our approaches significantly benefit this task compared to the PCA-based method, particularly in lower dimensions where the gap widens. The performance of the EmbedTextNet model improves at lower dimensions but still lags behind our WordFS-S methods. It is worth noting that even when reduced to 50 dimensions, our WordFS-S-wP method achieves better performance compared to the original 300-dimensional word embeddings. Our experiment with three different pre-trained word embeddings shows that our method achieved the best performance for the Glove and Fasttext embeddings when post-processing is applied. However, for the Word2vec embeddings, better results are obtained without post-processing. This is reasonable because the power of the post-processing method may depend on the pre-trained word embeddings and the specific application scenario. The Word2vec embeddings we used were pre-trained on the most extensive data among the three. It implies that the embeddings may be more resilient to noise and capable of capturing more nuanced information. However, applying PPA can potentially result in the loss of information and cause a drop in performance.

Table 2: Performance comparison of Spearman’s rank correlation coefficients on the aggregated word similarity dataset.

| Dimension | 300 | 150 | 100 | 50 |
|--------------------|-------|--------------|--------------|--------------|
| Glove | 45.74 | – | 38.45 | 36.24 |
| UMAP | – | 13.97 | 13.82 | 13.48 |
| Algo | – | 53.48 | 49.35 | 42.61 |
| EmbedTextNet | – | 52.71 | 43.88 | 41.00 |
| WordFS-P-woP(ours) | – | 48.23 | 48.39 | 48.22 |
| WordFS-P-wP(ours) | – | 54.69 | 52.60 | 47.29 |
| WordFS-S-woP(ours) | – | <u>55.85</u> | <u>54.45</u> | <u>49.32</u> |
| WordFS-S-wP(ours) | – | 56.76 | 54.73 | 50.11 |
| Word2vec | 59.51 | – | – | – |
| UMAP | – | 32.54 | 32.20 | 32.16 |
| Algo | – | 59.11 | 56.98 | 50.74 |
| EmbedTextNet | – | 58.81 | 57.53 | <u>54.15</u> |
| WordFS-P-woP(ours) | – | 59.06 | 57.37 | 53.58 |
| WordFS-P-wP(ours) | – | 59.15 | 57.20 | 52.99 |
| WordFS-S-woP(ours) | – | <u>59.53</u> | <u>58.04</u> | 54.82 |
| WordFS-S-wP(ours) | – | 59.76 | 58.26 | 54.13 |
| Fasttext | 54.63 | – | – | – |
| UMAP | – | 38.48 | 38.54 | 38.40 |
| Algo | – | 58.24 | 56.12 | 49.76 |
| EmbedTextNet | – | 56.78 | 55.11 | 51.17 |
| WordFS-P-woP(ours) | – | 52.16 | 46.55 | 35.49 |
| WordFS-P-wP(ours) | – | 60.86 | 57.69 | 50.46 |
| WordFS-S-woP(ours) | – | <u>64.22</u> | <u>61.65</u> | <u>56.87</u> |
| WordFS-S-wP(ours) | – | 65.18 | 63.54 | 57.43 |

4.3 Downstream Tasks

It is insufficient to only evaluate word embeddings based on word similarity tasks [7]. To demonstrate the power of our method, generalizability on downstream tasks is crucial. We utilize the SentEval toolkit [3] to evaluate the dimension-reduced vectors on various downstream tasks. We conducted experiments on nine prediction tasks and five sentence similarity tasks. The sentence representations were obtained by averaging the word embeddings. We utilize the WordFS-P for the prediction tasks and the WordFS-S method for the sentence similarity tasks.

Prediction tasks: We conduct experiments on nine prediction tasks from the SentEval toolkit, including binary and multi-class classification tasks (MR, CR, MPQA, SUBJ, STS-B, STS-FG, and TREC), paraphrase detection task (MRPC), and entailment and semantic relatedness task (SICK-E). These tasks directly take the word embeddings as input. We applied our method with the RFT feature selection based on weak guidance from the aggregated word similarity dataset to the pre-trained word embeddings in all the experiments to reduce the dimension. Then, we directly input the dimension-reduced embeddings to the downstream tasks and compare the test accuracy of different methods for each task.

Table 3 shows the results of various methods on the nine prediction tasks. The last column is the average of the previous columns. Figure 2 compares the average accuracy among the top three methods. Experiments show that our method surpasses the UMAP method in all the prediction tasks and outperforms the PCA-based method in most tasks. Also, our method achieves higher average scores than the PCA-based method in all the settings, except for a slightly lower performance when reducing the Fasttext word embedding from 300 to 50 dimensions, with a difference of only 0.02. Our method also performs better than the EmbedTextNet model, although it performs slightly worse when reduced to 50 dimensions. However, as shown in the next section, our model is much more efficient and takes less time to train than the EmbedTextNet model. Our method retains more helpful information for most settings while reducing dimensions, achieving a much closer prediction performance to the original 300-dimensional word embeddings. It indicates the effectiveness of our method. Also, it proves that our method can generalize well on various prediction downstream tasks in a zero-shot manner. The Glove word embedding we used was pre-trained on a smaller amount of data and might be vulnerable to noise. Post-processing can help mitigate noise and improve the performance. However, subtle linguistic connections can be crucial when working with prediction tasks. Applying PPA to word embeddings pre-trained on relatively large data may decrease performance.




Figure 2: Average accuracy comparison for prediction tasks.

Reducing the dimension of data results in information loss, often leading to decreased performance as the dimension decreases. Figure 3 compares the impact of word vector dimensions on performance across different tasks. The impact of dimension reduction on task performance can vary depending on the pre-trained word embeddings used, but the overall trend remains consistent. Various tasks may lose crucial information at different reduction levels. For example, the TREC task’s performance significantly drops when reduced to 50D compared to 100D, while the performance of most other tasks decreases gradually when reduced to a lower amount. When using Word2vec and Fasttext word embeddings, the TREC task’s performance

Table 3: Comparison of the results of downstream prediction tasks, where the last column is the average of the previous columns.

| Task | MR | CR | MPQA | SUBJ | STS-B | SST -FG | TREC | MRPC | SICK-E | Avg |
|-------------------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|--------------|--------------|--------------|
| Glove-300D | 74.99 | 75.81 | 86.75 | 91.39 | 78.20 | 40.77 | 66.6 | 72.58 | 77.19 | 73.81 |
| UMAP-150D | 50.00 | 63.76 | 68.77 | 50.00 | 49.92 | 28.64 | 18.8 | 67.48 | 56.69 | 50.45 |
| Algo-150D | <u>73.70</u> | <u>75.32</u> | 85.31 | 89.59 | 76.66 | <u>40.95</u> | 60.7 | 71.68 | <u>76.49</u> | 72.27 |
| EmbedTextNet-150D | 72.67 | 74.09 | 84.95 | 90.08 | 73.75 | 40.50 | 64.8 | 71.54 | 74.43 | 71.87 |
| WordFS-P-woP-150D(ours) | 73.09 | 74.86 | 86.24 | <u>89.86</u> | 76.22 | 39.55 | 65.4 | 71.01 | 76.62 | <u>72.54</u> |
| WordFS-P-wP-150D(ours) | 72.68 | 75.42 | 86.24 | 89.20 | <u>76.50</u> | 41.27 | <u>65.2</u> | 70.38 | 76.46 | 72.81 |
| UMAP-100D | 50.00 | 63.82 | 68.77 | 50.00 | 49.92 | 28.64 | 18.8 | 67.36 | 56.69 | 50.44 |
| Algo-100D | 70.61 | 72.82 | 82.71 | 88.16 | <u>74.63</u> | 38.28 | 56.2 | 71.83 | <u>75.36</u> | 70.07 |
| EmbedTextNet-100D | 71.55 | 72.87 | 84.44 | 88.88 | 70.18 | <u>38.46</u> | <u>62.2</u> | 71.48 | 75.14 | 70.58 |
| WordFS-P-woP-100D(ours) | <u>72.24</u> | <u>73.96</u> | <u>85.26</u> | <u>88.62</u> | 73.97 | 38.37 | 62.0 | 71.42 | 76.60 | <u>71.38</u> |
| WordFS-P-wP-100D(ours) | 72.68 | 75.42 | 85.50 | 88.25 | 76.94 | 40.50 | 64.4 | <u>71.59</u> | 74.87 | 72.24 |
| UMAP-50D | 50.00 | 63.76 | 68.78 | 50.00 | 49.92 | 28.64 | 18.8 | 67.83 | 56.69 | 50.49 |
| Algo-50D | 66.24 | 70.23 | 76.35 | 84.32 | 69.36 | 35.88 | 48.1 | <u>71.39</u> | 72.98 | 66.09 |
| EmbedTextNet-50D | <u>69.36</u> | 72.34 | <u>82.79</u> | 87.65 | <u>71.66</u> | 38.64 | 60.2 | 68.29 | 74.20 | 69.45 |
| WordFS-P-woP-50D(ours) | 68.54 | 70.54 | 81.35 | <u>85.51</u> | 69.85 | 35.61 | 56.2 | 71.65 | <u>73.19</u> | 68.05 |
| WordFS-P-wP-50D(ours) | 69.96 | <u>71.05</u> | 83.29 | <u>84.25</u> | 72.10 | <u>37.69</u> | 56.6 | 71.30 | <u>72.82</u> | <u>68.78</u> |
| Word2vec-300D | 77.16 | 77.80 | 87.97 | 90.42 | 81.11 | 42.22 | 82.6 | 72.35 | 77.92 | 76.62 |
| UMAP-150D | 50.03 | 63.76 | 68.74 | 50.00 | 49.92 | 28.64 | 18.8 | 66.49 | 56.69 | 50.34 |
| Algo-150D | <u>74.52</u> | 75.28 | 85.72 | <u>89.31</u> | <u>77.70</u> | 39.46 | 70.0 | 70.55 | 73.13 | 72.85 |
| EmbedTextNet-150D | 73.87 | <u>77.14</u> | 84.98 | 89.73 | 77.05 | 40.72 | <u>76.0</u> | 73.28 | 73.88 | 74.07 |
| WordFS-P-woP-150D(ours) | 74.43 | 76.90 | 86.56 | 88.72 | 79.13 | <u>40.18</u> | 77.0 | 72.12 | 76.54 | 74.65 |
| WordFS-P-wP-150D(ours) | 75.51 | 77.41 | <u>86.15</u> | 88.45 | 76.88 | 40.14 | 73.8 | <u>72.17</u> | 76.54 | <u>74.12</u> |
| UMAP-100D | 50.00 | 63.76 | 68.77 | 50.00 | 53.71 | 28.64 | 18.8 | 66.49 | 58.78 | 50.99 |
| Algo-100D | 72.10 | 74.70 | 84.42 | <u>88.30</u> | <u>76.11</u> | 38.60 | 62.2 | 72.87 | <u>74.08</u> | 71.49 |
| EmbedTextNet-100D | 72.93 | 76.40 | 85.43 | 89.33 | 77.81 | 36.65 | 65.8 | 72.29 | 73.78 | 72.27 |
| WordFS-P-woP-100D(ours) | <u>72.75</u> | 74.25 | <u>84.77</u> | 87.77 | 76.06 | <u>39.77</u> | 75.2 | <u>72.70</u> | 73.68 | 72.97 |
| WordFS-P-wP-100D(ours) | <u>72.75</u> | <u>74.83</u> | 84.65 | 86.52 | 75.78 | 40.05 | <u>73.0</u> | 70.09 | 75.24 | <u>72.54</u> |
| UMAP-50D | 50.00 | 63.76 | 68.70 | 50.00 | 52.50 | 28.64 | 13.00 | 66.49 | 56.69 | 49.98 |
| Algo-50D | <u>70.39</u> | 70.84 | <u>83.39</u> | <u>86.73</u> | <u>73.64</u> | 37.06 | 54.4 | 72.70 | 69.60 | 68.75 |
| EmbedTextNet-50D | 70.79 | 72.95 | 84.35 | 87.47 | 74.74 | 40.00 | 66.8 | 70.84 | 72.21 | 71.13 |
| WordFS-P-woP-50D(ours) | 69.39 | <u>71.60</u> | 81.35 | 85.54 | 72.71 | 37.69 | <u>62.6</u> | <u>72.52</u> | 70.90 | <u>69.37</u> |
| WordFS-P-wP-50D(ours) | 70.02 | 71.52 | 82.48 | 82.66 | 72.43 | <u>38.10</u> | 61.6 | 71.65 | <u>71.10</u> | 69.06 |
| Fasttext-300D | 77.59 | 78.81 | 88.10 | 91.73 | 80.78 | 45.48 | 84.4 | 72.99 | 77.57 | 77.49 |
| UMAP-150D | 50.00 | 63.76 | 68.80 | 50.00 | 49.92 | 28.64 | 18.8 | 71.07 | 56.69 | 50.85 |
| Algo-150D | 74.81 | 75.39 | 85.72 | 90.12 | 77.81 | 41.58 | 66.4 | 73.39 | 70.79 | 72.89 |
| EmbedTextNet-150D | 73.95 | 76.77 | 86.32 | <u>90.80</u> | 78.91 | 40.54 | <u>73.2</u> | 72.12 | <u>74.79</u> | 74.16 |
| WordFS-P-woP-150D(ours) | <u>75.68</u> | <u>75.89</u> | 86.32 | 89.62 | 79.68 | <u>42.04</u> | 76.0 | <u>72.29</u> | 74.24 | 74.64 |
| WordFS-P-wP-150D(ours) | 76.01 | 75.87 | 86.13 | 89.39 | <u>79.57</u> | 42.08 | 73.0 | 72.12 | 75.28 | <u>74.38</u> |
| UMAP-100D | 50.00 | 63.76 | 68.77 | 50.00 | 49.92 | 28.64 | 16.2 | 66.49 | 56.69 | 50.05 |
| Algo-100D | 73.31 | <u>76.05</u> | <u>85.43</u> | <u>88.96</u> | <u>77.27</u> | 40.50 | 66.0 | <u>72.58</u> | 70.98 | 72.34 |
| EmbedTextNet-100D | 74.56 | 76.74 | 86.17 | 90.09 | 77.87 | 41.22 | 74.8 | 72.52 | 74.41 | 74.26 |
| WordFS-P-woP-100D(ours) | <u>74.43</u> | 74.86 | 84.16 | 88.54 | 76.94 | 40.23 | <u>69.8</u> | 72.06 | <u>71.79</u> | <u>72.53</u> |
| WordFS-P-wP-100D(ours) | 73.51 | 73.56 | 85.06 | 86.33 | 76.83 | <u>40.54</u> | 69.4 | 73.04 | 71.34 | 72.18 |
| UMAP-50D | 50.04 | 63.76 | 68.79 | 50.00 | 49.92 | 28.64 | 32.6 | 66.90 | 56.69 | 51.93 |
| Algo-50D | 70.99 | <u>71.55</u> | <u>84.12</u> | <u>86.14</u> | <u>73.09</u> | 38.14 | 55.0 | <u>72.64</u> | 70.47 | <u>69.13</u> |
| EmbedTextNet-50D | 72.34 | 72.08 | 84.85 | 87.97 | 72.60 | 37.65 | 72.6 | 73.16 | 71.40 | 71.63 |
| WordFS-P-woP-50D(ours) | <u>71.82</u> | 70.33 | 81.43 | 83.98 | 74.63 | <u>37.87</u> | <u>60.0</u> | 71.71 | 70.20 | 69.11 |
| WordFS-P-wP-50D(ours) | 70.23 | 69.11 | 81.41 | 81.79 | 71.77 | 36.61 | 57.0 | 71.94 | <u>70.51</u> | 67.82 |



Figure 3: Accuracy comparison for prediction tasks using our WordFS method.

drops significantly when the dimension is reduced to 150D. This suggests that important information about the type of question is lost early on. For most other tasks, performance decreases slightly as the dimension decreases, which is expected when dimensions are reduced. In conclusion, dimension reduction models can preserve essential information for improved overall performance, but the performance of specific tasks may vary without further fine-tuning.

Similarity tasks: Furthermore, we evaluate our method on five sentence similarity tasks from the SentEval toolkit. The sentence similarity tasks contain STS tasks from 2012 to 2016, aiming to determine how close the distance between two sentence vectors correlates with a similarity score assigned by human annotators. We first applied our Sim-based method to reduce the word embeddings based on the weak guidance provided by the aggregated word similarity dataset. Then, we directly input the dimension-reduced vectors and compare the average Spearman’s rank correlation coefficients among different methods.

Table 4 shows the results of various methods on the sentence similarity tasks. Figure 4 compares the average Spearman’s rank correlation coefficients among different methods. Our method is consistently better than the existing methods, except when we reduce Glove embedding to 150 dimensions and Word2vec embedding to 50 dimensions, where the existing methods perform better with less than 0.05 difference. Our method significantly outperforms the existing approach in most sentence similarity tasks, with a gap of up to 8.36. The observation of post-processing is the same as that of the aggregated word similarity dataset. The Word2vec embeddings we use are trained on extensive data and may not benefit much from the PPA. The performance of the reduced word embeddings is often as good or even better than the original word vectors for Glove and Fasttext word embeddings in sentence similarity tasks. Even for Word2vec word embeddings, the drop in performance is small. This indicates that our method, which is based on word similarity, effectively retains critical dimensions for sentence similarity tasks while reducing the confusing ones.

Table 4: Performance comparison of downstream sentence similarity tasks, where the last column is the average of the previous columns.

| Task | STS -12 | STS -13 | STS -14 | STS -15 | STS -16 | Avg |
|-------------------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Glove-300D | 51.52 | 48.40 | 54.21 | 57.09 | 55.28 | 53.30 |
| UMAP-150D | 39.10 | 26.16 | 28.96 | 33.65 | 40.94 | 33.76 |
| Algo-150D | 53.23 | <u>56.41</u> | 62.30 | 67.52 | 66.99 | 61.29 |
| EmbedTextNet-150D | 52.68 | 52.12 | 58.67 | 61.87 | 59.80 | 57.03 |
| WordFS-S-woP-150D(ours) | 51.40 | 53.75 | 59.19 | 63.07 | 62.16 | 57.91 |
| WordFS-S-wP-150D(ours) | <u>52.74</u> | 58.62 | <u>62.16</u> | <u>66.42</u> | <u>66.24</u> | <u>61.24</u> |
| UMAP-100D | 39.69 | 26.90 | 28.97 | 33.98 | 41.40 | 26.79 |
| Algo-100D | 52.51 | <u>54.66</u> | 60.82 | 66.79 | 65.15 | 59.99 |
| EmbedTextNet-100D | 51.89 | 48.30 | 55.51 | 58.43 | 55.66 | 53.96 |
| WordFS-S-woP-100D(ours) | 49.83 | 51.00 | 57.98 | 61.59 | 62.74 | 56.63 |
| WordFS-S-wP-100D(ours) | <u>52.16</u> | 59.05 | 62.24 | 66.84 | 66.21 | 61.30 |
| UMAP-50D | 38.19 | 26.92 | 29.17 | 33.52 | 39.94 | 33.55 |
| Algo-50D | 48.20 | 49.44 | <u>57.53</u> | <u>61.49</u> | <u>61.42</u> | <u>55.62</u> |
| EmbedTextNet-50D | 48.83 | 45.86 | 53.57 | 56.80 | 53.32 | 51.68 |
| WordFS-S-woP-50D(ours) | 50.27 | <u>49.58</u> | 55.66 | 56.96 | 60.07 | 54.51 |
| WordFS-S-wP-50D(ours) | <u>49.71</u> | 55.42 | 60.58 | 63.50 | 63.62 | 58.57 |
| Word2vec-300D | 55.48 | 58.23 | 64.05 | 67.97 | 66.29 | 62.40 |
| UMAP-150D | 38.01 | 33.32 | 37.75 | 41.73 | 46.29 | 39.72 |
| Algo-150D | <u>54.94</u> | 58.63 | 64.01 | 65.90 | 65.14 | 61.72 |
| EmbedTextNet-150D | 54.85 | <u>58.91</u> | 64.31 | 66.79 | 65.45 | 62.06 |
| WordFS-S-woP-150D(ours) | 54.91 | 59.02 | 64.07 | 67.70 | <u>66.68</u> | 62.48 |
| WordFS-S-wP-150D(ours) | 55.12 | 58.90 | <u>64.30</u> | <u>67.13</u> | 66.73 | <u>62.44</u> |
| UMAP-100D | 37.87 | 35.87 | 38.15 | 42.11 | 46.92 | 40.18 |
| Algo-100D | 54.10 | 57.43 | <u>63.54</u> | 64.58 | 63.31 | 60.59 |
| EmbedTextNet-100D | 54.94 | 57.68 | 63.79 | 65.01 | 64.96 | 61.28 |
| WordFS-S-woP-100D(ours) | 55.38 | <u>57.77</u> | 63.32 | 67.20 | 66.27 | 61.99 |
| WordFS-S-wP-100D(ours) | <u>54.96</u> | 58.33 | 63.36 | <u>67.04</u> | <u>65.97</u> | <u>61.93</u> |
| UMAP-50D | 38.12 | 31.41 | 37.00 | 41.37 | 46.75 | 38.93 |
| Algo-50D | 51.20 | 53.46 | 60.51 | 61.89 | 60.76 | 57.56 |
| EmbedTextNet-50D | 54.23 | 55.85 | <u>61.81</u> | 62.64 | 62.69 | 59.44 |
| WordFS-S-woP-50D(ours) | 52.70 | 55.45 | 61.55 | 64.14 | 62.93 | 59.35 |
| WordFS-S-wP-50D(ours) | <u>53.03</u> | <u>55.61</u> | 61.99 | <u>63.64</u> | <u>62.83</u> | <u>59.42</u> |
| Fasttext-300D | 56.15 | 51.34 | 59.11 | 63.21 | 61.03 | 58.17 |
| UMAP-150D | 42.52 | 29.34 | 36.48 | 37.06 | 44.37 | 37.95 |
| Algo-150D | 56.80 | 55.03 | 61.72 | 64.07 | 61.05 | 59.73 |
| EmbedTextNet-150D | 55.22 | 51.03 | 59.55 | 62.67 | 60.94 | 57.88 |
| WordFS-S-woP-150D(ours) | <u>59.05</u> | <u>56.76</u> | <u>62.73</u> | <u>67.33</u> | <u>64.63</u> | <u>62.10</u> |
| WordFS-S-wP-150D(ours) | 60.72 | 62.79 | 64.82 | 69.95 | 70.02 | 65.66 |
| UMAP-100D | 41.49 | 28.76 | 36.17 | 37.09 | 43.99 | 37.50 |
| Algo-100D | 56.57 | <u>54.43</u> | 60.06 | 62.53 | 60.09 | 58.74 |
| EmbedTextNet-100D | 55.55 | 48.91 | 58.48 | 61.42 | 58.86 | 56.64 |
| WordFS-S-woP-100D(ours) | <u>57.74</u> | 53.78 | <u>61.07</u> | <u>65.60</u> | <u>62.37</u> | <u>60.11</u> |
| WordFS-S-wP-100D(ours) | 60.25 | 60.84 | 64.21 | 69.54 | 69.14 | 64.80 |
| UMAP-50D | 42.25 | 28.16 | 36.01 | 37.20 | 44.43 | 37.61 |
| Algo-50D | <u>55.56</u> | <u>53.82</u> | 59.99 | 60.60 | 68.91 | 57.78 |
| EmbedTextNet-50D | 55.34 | 49.76 | 58.98 | 59.70 | 58.52 | 56.46 |
| WordFS-S-woP-50D(ours) | 54.88 | 50.91 | <u>60.44</u> | <u>63.60</u> | 60.74 | <u>58.11</u> |
| WordFS-S-wP-50D(ours) | 57.33 | 56.74 | 60.94 | 65.41 | <u>64.10</u> | 60.90 |




Figure 4: Comparison of the average Spearman's rank correlation coefficients of sentence similarity tasks.




Figure 5: Spearman's rank correlation coefficients comparison for similarity tasks using our WordFS method.

Figure 5 compares the impact of word vector dimensions on performance across different tasks. As the dimensions decrease, there is generally a slight drop in performance because some detailed information is discarded. Most lines show a similar trend from 150D to 50D. For the Glove and Word2vec embeddings, our approach significantly improves performance when reducing to 150D. This is understandable because these embeddings are trained on a smaller amount of data, and our approach helps the vectors focus on similarity tasks. Word2vec word embeddings are trained on a much larger amount of data, which is sufficient for them to perform well on sentence similarity tasks.

In conclusion, our method generally outperforms existing methods in prediction and similarity tasks, demonstrating its superiority and generalizability.

4.4 Model Efficiency

We compare the model efficiency of the Algo method, the EmbedTextNet method, and our method when reducing the word embedding dimensions from 300 to 150 in Table 5. For each pre-trained word embedding, all methods start with the same vocabulary of 300D word vectors and output a file with 150D word vectors. Different hardware is used for other methods because of

Table 5: Complexity comparison of the Algo method, the EmbedTextNet method, and our methods when reducing the word embedding dimensions from 300 to 150 in hardware configuration and training time.

| Embeddings | Method | Hardware Configuration | Training Time (s) |
|------------|---------------------|------------------------|----------------------|
| Glove | EmbedTextNet | GPU | 2149.04 (5372.60X) |
| | UMAP | CPU | 655.52 (1638.80X) |
| | Algo | | 21.50 (53.75X) |
| | WordFS-S-wP (ours) | | 17.63 (44.08X) |
| | WordFS-P-wP (ours) | | 11.84 (29.60X) |
| | WordFS-S-woP (ours) | | 4.05 (10.13X) |
| | WordFS-P-woP (ours) | | 0.40 (1X) |
| Word2vec | EmbedTextNet | GPU | 18272.62 (30970.54X) |
| | UMAP | CPU | 8470.38 (14356.58X) |
| | Algo | | 198.93 (337.17X) |
| | WordFS-S-wP (ours) | | 93.88 (159.12X) |
| | WordFS-P-wP (ours) | | 85.93 (145.64X) |
| | WordFS-S-woP (ours) | | 7.36 (12.47X) |
| | WordFS-P-woP (ours) | | 0.59 (1X) |
| Fasttext | EmbedTextNet | GPU | 5306.65 (3537.77X) |
| | UMAP | CPU | 2971.88 (1982.59X) |
| | Algo | | 74.98 (49.99X) |
| | WordFS-S-wP (ours) | | 52.20 (34.80X) |
| | WordFS-P-wP (ours) | | 36.54 (24.36X) |
| | WordFS-S-woP (ours) | | 15.24 (10.16X) |
| | WordFS-P-woP (ours) | | 1.50 (1X) |

their algorithmic nature. All the experiments were conducted on the same server, equipped with two AMD EPYC 7543 32-core CPUs and seven NVIDIA RTX A6000 GPUs. The deep learning-based model is trained using the GPUs, while the Algo model and our methods are trained using only the CPUs. The EmbedTextNet model takes approximately 30 minutes to train on relatively small GloVe word embeddings.

In contrast, the training time for our WordFS-P-wP and WordFS-S-wP models is less than 19 seconds and 13 seconds, respectively. Our methods without post-processing require even less time — the training of the WordFS-P-woP model is completed within one second. For more significant pre-trained word embeddings, it takes hours to train the EmbedTextNet model. However, our WordFS-P-wP and WordFS-S-wP methods only require 1-2 minutes, making them hundreds of times faster. Our WordFS-P-woP model, without post-processing, consistently completes in around one second, making it thousands or ten thousand times faster than the deep learning-based method. Besides, our method is significantly faster and outperforms UMAP. Because finding the K nearest neighbors may take a long time in a high-dimensional word vector space.

Additionally, our methods are more efficient than the PCA-based method because they employ PPA before and after the PCA, while we only apply it once before feature selection. The PPA consumes a substantial portion of the training time compared to the feature selection module. Our model, without the post-processing, can achieve even less time. The reason is the post-processing and PCA must be done on the entire vocabulary. However, our feature selection method only focuses on a subset of words that appear in the aggregated word similarity dataset, making the procedure very fast. And the module’s processing time will not be significantly affected by the vocabulary size.

In conclusion, Our methods take only a fraction of the time compared to deep learning methods and less than half the time compared to existing PCA-based methods. The experimental results demonstrate that our methods are much more efficient than existing methods.

5 Conclusion and Future Work

This paper proposes a general dimension reduction method called WordFS for pre-trained word embeddings with the post-processing algorithm (PPA) and simple yet effective feature selection methods based on weak supervision provided by a limited number of word similarity pairs. Our method is more straightforward, efficient, intuitive, and effective in most cases than the existing methods. Empirical results show that our method outperforms existing methods in word similarity tasks and generalizes well to various downstream tasks. Our model demonstrates clear advantages for tasks significantly related to word similarity, like STS tasks. Even for tasks that might not directly correlate with word similarity, such as classification tasks, our model performs better on average, showing our approach’s generalizability. We demonstrate that our proposed weakly-supervised feature selection method can effectively reduce word embedding dimensions and generalize to many downstream tasks with much lower computational costs. In the future, we would like to explore the compression of word embeddings further. Since feature selection methods may result in information loss, there is still room for improvement in the performance of prediction tasks between the original word embeddings and the reduced word vectors. Additionally, suitable feature selection methods can be developed to compress pre-trained word vectors in specific domains.

References

- [1] M. Ait-Saada and M. Nadif, “Is anisotropy truly harmful? a case study on text clustering”, in *Proceedings of the 61st Annual Meeting of the*

- Association for Computational Linguistics (Volume 2: Short Papers)*, 2023, 1194–203.
- [2] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, “Enriching Word Vectors with Subword Information”, *arXiv preprint arXiv:1607.04606*, 2016.
 - [3] A. Conneau and D. Kiela, “SentEval: An Evaluation Toolkit for Universal Sentence Representations”, *arXiv preprint arXiv:1803.05449*, 2018.
 - [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding”, *arXiv preprint arXiv:1810.04805*, 2018.
 - [5] Z. Erenel, O. R. Adegboye, and H. Kusetogullari, “A new feature selection scheme for emotion recognition from text”, *Applied Sciences*, 10(15), 2020, 5351.
 - [6] M. Faruqui and C. Dyer, “Improving vector space word representations using multilingual correlation”, in *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 2014, 462–71.
 - [7] M. Faruqui, Y. Tsvetkov, P. Rastogi, and C. Dyer, “Problems with evaluation of word embeddings using word similarity tasks”, *arXiv preprint arXiv:1605.02276*, 2016.
 - [8] D. Y. Hwang, B. Taha, and Y. Nechaev, “Embedtextnet: Dimension reduction with weighted reconstruction and correlation losses for efficient text embedding”, in *Findings of the Association for Computational Linguistics: ACL 2023*, 2023, 9863–79.
 - [9] R. Jha and K. Mihata, “On geodesic distances and contextual embedding compression for text classification”, *arXiv preprint arXiv:2104.11295*, 2021.
 - [10] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu, “Tinybert: Distilling bert for natural language understanding”, *arXiv preprint arXiv:1909.10351*, 2019.
 - [11] Y. Kim, K.-M. Kim, and S. Lee, “Adaptive compression of word embeddings”, in *Proceedings of the 58th annual meeting of the association for computational linguistics*, 2020, 3950–9.
 - [12] C.-C. J. Kuo and A. M. Madni, “Green learning: Introduction, examples and outlook”, *Journal of Visual Communication and Image Representation*, 90, 2023, 103685.
 - [13] L. McInnes, J. Healy, and J. Melville, “Umap: Uniform manifold approximation and projection for dimension reduction”, *arXiv preprint arXiv:1802.03426*, 2018.
 - [14] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space”, *arXiv preprint arXiv:1301.3781*, 2013.

- [15] T. Mikolov, E. Grave, P. Bojanowski, C. Puhersch, and A. Joulin, “Advances in Pre-Training Distributed Word Representations”, in *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [16] J. Mu and P. Viswanath, “All-but-the-Top: Simple and Effective Post-processing for Word Representations”, in *International Conference on Learning Representations*, 2018, <https://openreview.net/forum?id=HkuGJ3kCb>.
- [17] H. Naveed, A. U. Khan, S. Qiu, M. Saqib, S. Anwar, M. Usman, N. Barnes, and A. Mian, “A comprehensive overview of large language models”, *arXiv preprint arXiv:2307.06435*, 2023.
- [18] J. Pennington, R. Socher, and C. D. Manning, “Glove: Global vectors for word representation”, in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, 1532–43.
- [19] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, “Deep contextualized word representations”, 2018, arXiv: [1802.05365](https://arxiv.org/abs/1802.05365) [cs.CL].
- [20] V. Raunak, V. Gupta, and F. Metze, “Effective dimensionality reduction for word embeddings”, in *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, 2019, 235–43.
- [21] W. Rui, J. Liu, and Y. Jia, “Unsupervised feature selection for text classification via word embedding”, in *2016 IEEE International Conference on Big Data Analysis (ICBDA)*, IEEE, 2016, 1–5.
- [22] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter”, *arXiv preprint arXiv:1910.01108*, 2019.
- [23] R. Sheikhpour, M. A. Sarram, S. Gharaghani, and M. A. Z. Chahooki, “A survey on semi-supervised feature selection methods”, *Pattern recognition*, 64, 2017, 141–58.
- [24] R. Shu and H. Nakayama, “Compressing word embeddings via deep compositional code learning”, *arXiv preprint arXiv:1711.01068*, 2017.
- [25] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction”, *science*, 290(5500), 2000, 2319–23.
- [26] A. K. Uysal and Y. L. Murphey, “Sentiment classification: Feature selection based approaches versus deep learning”, in *2017 IEEE International Conference on Computer and Information Technology (CIT)*, IEEE, 2017, 23–30.
- [27] L. Van der Maaten and G. Hinton, “Visualizing data using t-SNE.”, *Journal of machine learning research*, 9(11), 2008.
- [28] B. Wang, A. Wang, F. Chen, Y. Wang, and C.-C. J. Kuo, “Evaluating word embedding models: Methods and experimental results”, *APSIPA transactions on signal and information processing*, 8, 2019, e19.

- [29] C. Wei, Y.-C. Wang, B. Wang, and C.-C. J. Kuo, “An Overview of Language Models: Recent Developments and Outlook”, *APSIPA Transactions on Signal and Information Processing*, 13(2), 2024, DOI: [10.1561/116.00000010](https://doi.org/10.1561/116.00000010), <http://dx.doi.org/10.1561/116.00000010>.
- [30] T. Wu, S. He, J. Liu, S. Sun, K. Liu, Q.-L. Han, and Y. Tang, “A brief overview of ChatGPT: The history, status quo and potential future development”, *IEEE/CAA Journal of Automatica Sinica*, 10(5), 2023, 1122–36.
- [31] J. Xue, Y.-C. Wang, C. Wei, X. Liu, J. Woo, and C.-C. J. Kuo, “Bias and Fairness in Chatbots: An Overview”, *APSIPA Transactions on Signal and Information Processing*, 13(2), 2024, DOI: [10.1561/116.00000064](https://doi.org/10.1561/116.00000064), <http://dx.doi.org/10.1561/116.00000064>.
- [32] Y. Yang, W. Wang, H. Fu, and C.-C. J. Kuo, “On supervised feature selection from high dimensional feature spaces”, *APSIPA Transactions on Signal and Information Processing*, 11(1), 2022.
- [33] G. Zhang, Y. Zhou, and D. Bollegala, “Evaluating Unsupervised Dimensionality Reduction Methods for Pretrained Sentence Embeddings”, *arXiv preprint arXiv:2403.14001*, 2024.