

# Unsupervised Sentence-embeddings by Manifold Approximation and Projection

Subhradeep Kayal

Prosus N.V.

Amsterdam, Netherlands.

deep.kayal@pm.me

## Abstract

The concept of unsupervised *universal sentence encoders* has gained traction recently, wherein pre-trained models generate effective task-agnostic fixed-dimensional representations for phrases, sentences and paragraphs. Such methods are of varying complexity, from simple weighted-averages of word vectors to complex language-models based on bidirectional transformers. In this work we propose a novel technique to generate sentence-embeddings in an unsupervised fashion by projecting the sentences onto a fixed-dimensional manifold with the objective of preserving local neighbourhoods in the original space. To delineate such neighbourhoods we experiment with several set-distance metrics, including the recently proposed *Word Mover's distance*, while the fixed-dimensional projection is achieved by employing a scalable and efficient manifold approximation method rooted in topological data analysis. We test our approach, which we term *EMAP* or *Embeddings by Manifold Approximation and Projection*, on six publicly available text-classification datasets of varying size and complexity. Empirical results show that our method consistently performs similar to or better than several alternative state-of-the-art approaches.

## 1 Introduction

### 1.1 On sentence-embeddings

Dense vector representation of words, or *word-embeddings*, form the backbone of most modern NLP applications and can be constructed using context-free (Bengio et al., 2003; Mikolov et al., 2013; Pennington et al., 2014) or contextualized methods (Peters et al., 2018; Devlin et al., 2019).

Given that practical systems often benefit from having representations for sentences and documents, in addition to word-embeddings (Palangi et al., 2016; Yan et al., 2016), a simple trick is

to use the weighted average over some or all of the embeddings of words in a sentence or document. Although *sentence-embeddings* constructed this way often lose information because of the disregard for word-order during averaging, they have been found to be surprisingly performant (Aldarmaki and Diab, 2018).

More sophisticated methods focus on jointly learning the embeddings of sentences and words using models similar to *Word2Vec* (Le and Mikolov, 2014; Chen, 2017), using encoder-decoder approaches that reconstruct the surrounding sentences of an encoded passage (Kiros et al., 2015), or training bi-directional LSTM models on large external datasets (Conneau et al., 2017). Meaningful sentence-embeddings have also been constructed by fine-tuning pre-trained bidirectional transformers (Devlin et al., 2019) using a Siamese architecture (Reimers and Gurevych, 2019).

In parallel to the approaches mentioned above, a stream of methods have emerged recently which exploit the inherent geometric properties of the structure of sentences, by treating them as sets or sequences of word-embeddings. For example, Arora et al. (2017) propose the construction of sentence-embeddings based on weighted word-embedding averages with the removal of the dominant singular vector, while Rücklé et al. (2018) produce sentence-embeddings by concatenating several power-means of word-embeddings corresponding to a sentence. Very recently, spectral decomposition techniques were used to create sentence-embeddings, which produced state-of-the-art results when used in concatenation with averaging (Kayal and Tsatsaronis, 2019; Almarwani et al., 2019).

Our work is most related to that of Wu et al. (2018) who use Random Features (Rahimi and Recht, 2008) to learn document embeddings which preserve the properties of an explicitly-defined kernel based on the Word Mover's Distance (Kusner

et al., 2015). Where Wu et al. predefine the nature of the kernel, our proposed approach can learn the similarity-preserving manifold for a given set-distance metric, offering increased flexibility.

## 1.2 Motivation and contributions

A simple way to form sentence-embeddings is to compute the dimension-wise arithmetic mean of the embeddings of the words in a particular sentence. Even though this approach incurs information loss by disregarding the fact that sentences are sequences (or, at the very least, sets) of word vectors, it works well in practice. This already provides an indication that there is more information in the sentences to be exploited.

Kusner et al. (2015) aim to use more of the information available in a sentence by representing sentences as a weighted point cloud of embedded words. Rooted in transportation theory, their *Word Mover’s distance (WMD)* is the minimum amount of distance that the embedded words of a sentence need to travel to reach the embedded words of another sentence. The approach achieves state-of-the-art results for sentence classification when combined with a  $k$ -NN classifier (Cover and Hart, 1967). Since their work, other distance metrics have been suggested (Singh et al., 2019; Wang et al., 2019), also motivated by how transportation problems are solved.

Considering that sentences are sets of word vectors, a large variety of methods exist in literature that can be used to calculate the distance between two sets, in addition to the ones based on transport theory. Thus, as a *first contribution*, we compare alternative metrics to measure distances between sentences. The metrics we suggest, namely the *Hausdorff distance* and the *Energy distance*, are intuitive to explain and reasonably fast to calculate. The choice of these particular distances are motivated by their differing origins and their general usefulness in the respective application domains.

Once calculated, these distances can be used in conjunction with  $k$ -nearest neighbours for classification tasks, and  $k$ -means for clustering tasks. However, these learning algorithms are rather simplistic and the state-of-the-art machine learning algorithms require a fixed-length feature representation as input to them. Moreover, having fixed-length representations for sentences (*sentence-embeddings*) also provides a large degree of flexibility for downstream tasks, as compared to hav-

ing only relative distances between them. With this as motivation, the *second contribution* of this work is to produce sentence-embeddings that approximately preserve the topological properties of the original sentence space. We propose to do so using an efficient scalable manifold-learning algorithm termed *UMAP* (McInnes et al., 2018) from topological data analysis. Empirical results show that this process yields sentence-embeddings that deliver near state-of-the-art classification performance with a simple classifier.

## 2 Methodology

### 2.1 Calculating distances

In this work, we experiment with three different distance measures to determine the distance between sentences. The first measure (Energy distance) is motivated by a useful linkage criterion from hierarchical clustering (Rokach and Maimon, 2005), while the second one (Hausdorff distance) is an important metric from algebraic topology that has been successfully used in document indexing (Tsatsaronis et al., 2012). The final metric (Word Mover’s distance) is a recent extension of an existing distance measure between distributions, that is particularly suited for use with word-embeddings (Kusner et al., 2015).

Prior to defining the distances that have been used in this work, we first proceed to outline the notations that we will be using to describe them.

#### 2.1.1 Notations

Let  $\mathcal{W} \in \mathbb{R}^{N \times d}$  denote a word-embedding matrix, such that the vocabulary corresponding to it consists of  $N$  words, and each word in it,  $w_i \in \mathbb{R}^d$ , is  $d$ -dimensional. This word-embedding matrix and its constituent words may come from pre-trained representations such as Word2Vec (Mikolov et al., 2013) or GloVe (Pennington et al., 2014), in which case  $d = 300$ .

Let  $\mathcal{S}$  be a set of sentences and  $s, s'$  be two sentences from this set. Each such sentence can be viewed as a set of word-embeddings,  $\{w\} \in s$ . Additionally, let the length of a sentence,  $s$ , be denoted as  $|s|$ , and the cardinality of the set,  $\mathcal{S}$ , be denoted by  $|\mathcal{S}|$ .

Let  $e(w_i, w_j)$  denote the distance between two word-embeddings,  $w_i, w_j$ . In the context of this paper, this distance is Euclidean:

$$e(w_i, w_j) = \|w_i - w_j\|_2 \quad (1)$$

Finally,  $D(s, s')$  denotes the distance between two sentences.

### 2.1.2 Energy distance

*Energy distance* is a statistical distance between probability distributions, based on the inter and intra-distribution variance, that satisfies all the criteria of being a metric (Székely and Rizzo, 2013).

Using the notations defined earlier, we write it as:

$$D(s, s') = \frac{2}{|s||s'|} \sum_{w_i \in s} \sum_{w_j \in s'} e(w_i, w_j) - \frac{1}{|s|^2} \sum_{w_i \in s} \sum_{w_j \in s} e(w_i, w_j) - \frac{1}{|s'|^2} \sum_{w_i \in s'} \sum_{w_j \in s'} e(w_i, w_j) \quad (2)$$

The original conception of the energy distance was inspired by gravitational potential energy of celestial objects. Looking closely at Equation 2, it can be quickly observed that it has two parts: the first term resembles the attraction or repulsion between two objects (or in our case, sentences), while the second and the third term indicate the self-coherence of the respective objects. As shown by Székely and Rizzo (2013), energy distance is scale equivariant, which would make it sensitive to contextual changes in sentences, and therefore make it useful in NLP applications.

### 2.1.3 Hausdorff distance

Given two subsets of a metric space, the *Hausdorff distance* is the maximum distance of the points in one subset to the nearest point in the other. A significant work has gone into making it fast to calculate (Atallah, 1983) so that it can be applied to real-world problems, such as shape-matching in computer vision (Dubuisson and Jain, 1994).

To calculate it, the distance between each point from one set and the closest point from the other set is determined first. Then, the Hausdorff distance is calculated as the maximal point-wise distance. Considering sentences  $\{s, s'\}$  as subsets of word-embedding space,  $\mathbb{R}^{d \times N}$ , the *directed* Hausdorff distance can be given as:

$$h(s, s') = \max_{w_i \in s} \min_{w_j \in s'} e(w_i, w_j) \quad (3)$$

such that the symmetric Hausdorff distance is:

$$D(s, s') = \max\{h(s, s'), h(s', s)\} \quad (4)$$

### 2.1.4 Word Mover's distance

In addition to the representation of a sentence as a set of word-embeddings, a sentence  $s$  can also be represented as a  $N$ -dimensional normalized term-frequency vector, where  $n_i^s$  is the number of times word  $w_i$  occurs in sentence  $s$  normalized by the total number of words in  $s$ :

$$n_i^s = \frac{c_i^s}{\sum_{k=1}^N c_k^s} \quad (5)$$

where,  $c_i^s$  is the number of times word  $w_i$  appears in sentence  $s$ .

The goal of the *Word Mover's distance (WMD)* (Kusner et al., 2015) is to construct a sentence similarity metric based on the distances between the individual words within each sentence, given by Equation 1. In order to calculate the distance between two sentences, WMD introduces a transport matrix,  $T \in \mathbb{R}^{N \times N}$ , such that each element in it,  $T_{ij}$ , denotes how much of  $n_i^s$  should be transported to  $n_j^{s'}$ . Then, the WMD between two sentences is given as the solution of the following minimization problem:

$$D(s, s') = \min_{T \geq 0} \sum_{i,j=1}^N T_{ij} e(i, j) \quad \text{subject to, } \sum_{j=1}^N T_{ij} = n_i^s \text{ and } \sum_{i=1}^N T_{ij} = n_j^{s'} \quad (6)$$

Thus, WMD between two sentences is defined as the minimum distance required to transport the words from one sentence to another.

## 2.2 Generating neighbourhood-preserving embeddings via non-linear manifold-learning

In this work, we propose to construct sentence-embeddings which preserve the neighbourhood around sentences delineated by the relative distances between them. We posit that preserving the local neighbourhoods will serve as a proxy for preserving the original topological properties.

In order to learn a topology-preserving fixed-dimensional manifold, we seek inspiration from methods in non-linear dimensionality-reduction (Lee and Verleysen, 2007) and topological data analysis literature (Carlsson, 2009). When broadly categorized, these techniques consist of methods, such as *Locally Linear Embedding* (Roweis and Saul, 2000), that preserve local distances between

points, or those like *Stochastic Neighbour Embedding* (Hinton and Roweis, 2003; van der Maaten and Hinton, 2008) that preserve the conditional probabilities of points being neighbours. However, existing manifold-learning algorithms suffer from two shortcomings: they are computationally expensive and are often restricted in the number of output dimensions. In our work we use a method termed *Uniform Manifold Approximation and Projection* (UMAP) (McInnes et al., 2018), which is scalable and has no computational restrictions on the output embedding dimension.

The building block of UMAP is a particular type of a simplicial complex, known as the Vietoris-Rips complex. Recalling that a  $k$ -simplex is a  $k$ -dimensional polytope which is the convex hull of its  $k + 1$  vertices, and a simplicial complex is a set of simplices of various orders, the Vietoris-Rips simplicial complex is a collection of 0 and 1-simplices. In essence, this is a means to building a simple neighbourhood graph by connecting the original data points.




Figure 1: Figure showing a simple example of the embedding algorithm. On the left is the original sentence-space, approximated by the nearest neighbours graph formed by the Vietoris-Rips complex. Instead of points and edges, our simplicial complex has sets of points and edges between them, formed by one of the distance metrics mentioned in Section 2.1. In this example, four sentences, denoted by  $S1$  through  $S4$ , form two simplices, with  $S4$  being a 0-simplex. The sentences are denoted by colored ellipses, while the high-dimensional embedding of each word in a sentence is depicted by a point having the same color as the parent sentence ellipse. The UMAP algorithm is then employed to find a similarity-preserving Euclidean embedding-space, shown on the right, by minimizing the cross-entropy between the two representations.

A key difference, in this work, to the original

formulation is that an individual data sample (i.e., the vertex of a simplex) is not a  $d$ -dimensional point but a set of  $d$ -dimensional words that make up a sentence. By using any of the distance metrics defined in Section 2.1, it is possible to construct the simplicial complex that UMAP needs in order to build the topological representation of the original sentence space. An illustration can be found in Figure 1.

As per the formulation laid out for UMAP, the similarity between sentences  $s'$  and  $s$  is defined as:

$$v_{s'|s} = \exp \frac{-(D(s, s') - \rho_s)}{\sigma_s} \quad (7)$$

where  $\sigma_s$  is a normalisation factor selected based on an empirical heuristic (See Algorithm 3 in the work of McInnes et al. 2018),  $D(s, s')$  is the distance between two sentences as outlined by Equation 2, 4 or 6, and  $\rho_s$  is the distance of  $s$  from its nearest neighbour. It is worth mentioning that for scalability,  $v_{s'|s}$  is calculated only for predefined set of approximate nearest neighbours, which is a user-defined input parameter to the UMAP algorithm, using the efficient *nearest-neighbour descent* algorithm (Dong et al., 2011).

The similarity depicted in Equation 7 is asymmetric, and symmetrization is carried out by a fuzzy set union using the probabilistic t-conorm:

$$v_{ss'} = (v_{s'|s} + v_{s|s'}) - v_{s'|s}v_{s|s'} \quad (8)$$

As UMAP builds a Vietoris-Rips complex governed by Equation 7, it can take advantage of the nerve theorem (Borsuk, 1948), which makes this construction a homotope of the original topological space. In our case, this implies that we can build a simple nearest neighbours graph from a given corpus of sentences, which has certain guarantees of approximating the original topological space, as defined by the aforementioned distance metrics.

The next step is to define a similar nearest neighbours graph in a fixed low-dimensional Euclidean space. Let  $s_E, s'_E \in \mathbb{R}^{d_E}$  be the corresponding  $d_E$ -dimensional sentence-embeddings. Then the low dimensional similarities are given by:

$$w_{ss'} = (1 + a\|s_E - s'_E\|_2^b)^{-1} \quad (9)$$

where,  $\|s_E - s'_E\|$  is the Euclidean distance between the  $d_E$ -dimensional embeddings, and setting  $a, b$  are input-parameters, set to 1.929 and 0.791, respectively, as per the original implementation.



---

**Algorithm 1:** Constructing sentence-Embeddings by Manifold Approximation and Projection: *EMAP*

---

**Data:** A pre-trained word-embeddings matrix,  $\mathcal{W}$ ; a set of sentences,  $\mathcal{S}$ ; desired dimension of the generated sentence-embeddings,  $d_E$

**Result:** A set of sentence-embeddings,  $\{s_E\} \in \mathcal{S}_E$

- 1 Calculate the distance matrix for the entire set of sentences, such that the distance between any two sentences is given by Equation 2, 4 or 6;
  - 2 Using this distance matrix, calculate the nearest neighbour graph between all input sentences, given by Equations 7 and 8;
  - 3 Calculate the initial guess for the low dimensional embeddings,  $\mathcal{S}_E \in \mathbb{R}^{|\mathcal{S}| \times d_E}$ , using the graph laplacian of the original nearest neighbour graph;
  - 4 Until convergence, minimize the cross-entropy between the two representations (Equation 10) using stochastic gradient descent;
  - 5 Return the set of  $d_E$ -dimensional sentence-embeddings,  $\mathcal{S}_E$ ;
- 

The final step of the process is to optimize the low dimensional representation to have as close a fuzzy topological representation as possible to the original space. UMAP proceeds to do so by minimizing the cross-entropy between the two representations:

$$C = \sum_{s \neq s'} v_{ss'} \log \frac{v_{ss'}}{w_{ss'}} + (1 - v_{ss'}) \log \frac{1 - v_{ss'}}{1 - w_{ss'}} \quad (10)$$

usually done via stochastic gradient descent.

A summary of the proposed process used to produce sentence-embeddings is provided in Algorithm 1, and pictorially presented in Figure 1.

### 3 Datasets and resources

#### 3.1 Datasets

Six public datasets<sup>1</sup> have been used to empirically validate the method proposed in this paper. These datasets are of varying sizes, tasks and complexities, and have been used widely in existing liter-

<sup>1</sup>[https://drive.google.com/open?id=1sGgAo2SB0YKhQQK\\_kilUp8KSToCI55j1](https://drive.google.com/open?id=1sGgAo2SB0YKhQQK_kilUp8KSToCI55j1)

ature, thereby making comparisons and reporting possible. Information about the datasets can be found in Table 1.

#### 3.2 Resources

**Pre-trained word-embedding corpus:** We use the pre-trained set of word-embeddings provided by Mikolov et al (2013)<sup>2</sup>.

**Software implementations:** We use a variety of software packages and custom-written programs perform our experiments, the starting point being the calculation of sentence-wise distances. We calculate the Hausdorff distance using a directed implementation provided in the *Scipy* python library<sup>3</sup>, whereas the energy distance is calculated using *dcor*<sup>4</sup>. Lastly, the word mover’s distance is calculated using implementation provided by Kusner et al. (2015)<sup>5</sup>. In order to produce the symmetric distance matrix for a dataset, we employ custom parallel implementation which distributes the calculations over all available logical cores in a machine.

To calculate the sentence-embeddings, the implementation of UMAP provided by McInnes et al (2018) is used<sup>6</sup>. Finally, the classification is done via linear kernel support vector machines from the *scikit-learn* library (Pedregosa et al., 2011)<sup>7</sup>.

All of the code and datasets have been packaged and released<sup>8</sup> to rerun all of the experiments.

**Compute infrastructure:** All experiments were run on a *m4.2xlarge* machine on AWS-EC2<sup>9</sup>, which has 8 virtual CPUs and 32GB of RAM.

## 4 Experiments

### 4.1 Competing methods

In order to check the usefulness of our proposed approach, we benchmark its performance in two different ways. The first, and most obvious, approach is to consider the performance of the  $k$ -NN

<sup>2</sup><https://drive.google.com/file/d/0B7XkCwpI5KDYN1NUTTT1SS21pQmM/edit>

<sup>3</sup>[https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.distance.directed\\_hausdorff.html](https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.distance.directed_hausdorff.html)

<sup>4</sup>[https://dcor.readthedocs.io/en/latest/functions/dcor.energy\\_distance.html#dcor.energy\\_distance](https://dcor.readthedocs.io/en/latest/functions/dcor.energy_distance.html#dcor.energy_distance)

<sup>5</sup><https://github.com/mkusner/wmd>

<sup>6</sup><https://umap-learn.readthedocs.io/en/latest/api.html>

<sup>7</sup><https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVC.html>

<sup>8</sup><https://github.com/DeepK/distance-embed>

<sup>9</sup><https://aws.amazon.com/ec2/>

Dataset	#classes	#train docs	#test docs	#avg tokens	Data-details
amazon	4	5600	2400	70	Reviews labeled by product
bbsport	5	517	220	192	Articles labeled by sport
classic	4	4965	2128	62	Manuscripts labeled by publisher
ohsumed	10	3999	5153	104	Medical abstracts categorized by subject headings
reuters8	8	5485	2189	69	News article categorization
twitter	3	2176	932	8	Tweet sentiment analysis

Table 1: **Dataset information:** Metadata describing the datasets used in our experiments.

classifier as a baseline. This is motivated by the state-of-the-art  $k$ -NN based classification accuracy reported by Kusner et al. for the word mover’s distance. Thus, our embeddings need to match or surpass the performance of a  $k$ -NN based approach, in order to be considered for practical use.

The second approach is to compare the classification accuracies of several state-of-the-art embedding-generation algorithms on our chosen datasets. These are:

**dct** (Almarwani et al., 2019): embeddings are generated by employing discrete cosine transform on a set of word vectors.

**eigensent** (Kayal and Tsatsaronis, 2019): sentence representations produced via higher-order dynamic mode decomposition (Le Clainche and Vega, 2017) on a sequence of word vectors.

**wmovers** (Wu et al., 2018): a competing method which can learn sentence representations from the word mover’s distance based on kernel learning, termed in the original work as *word mover’s embeddings*.

**p-means** (Rücklé et al., 2018): produces sentence-embeddings by concatenating several power-means of word-embeddings corresponding to a sentence.

**doc2vec** (Le and Mikolov, 2014): embeddings produced by jointly learning the representations of sentences, together with words, as a part of the word2vec procedure.

**s-bert** (Reimers and Gurevych, 2019): embeddings produced by fine-tuning a pre-trained BERT model using a Siamese architecture to classify two sentences as being similar or different.

Note that the results for *wmovers* and *doc2vec* are taken from Table 3 of Wu et al.’s work (2018), while all the other algorithms are explicitly tested.

## 4.2 Setup

Extensive experiments are performed to provide a holistic overview of our neighbourhood-preserving embedding algorithm, for various sets of input parameters. The steps involved are as follows:

**Choose a dataset** (one of the six mentioned in

Section 3.1). For every word in every sentence in the train and test splits of the dataset, retrieve the corresponding word-embedding from the pre-trained embedding corpus (as stated in Section 3.2).

**Calculate symmetric distance matrices** corresponding to each of the chosen distance metrics, for all of the sets of word-embeddings from the train and test splits.

**Apply the UMAP algorithm** on the distance matrices to generate embeddings for all sentences in the train and the test splits.

**Calculate embeddings for competing methods** for the methods outlined in Section 4.1.

Embeddings are generated for various hyperparameter combinations for *EMAP* as well as all the compared approaches, as listed in Table 2.

**Train a classifier on the produced embeddings** to perform the dataset-specific task. In this work, we train a simple linear-kernel support vector machine (Cortes and Vapnik, 1995) for every competing method and every dataset tested. The classifier is trained on the train-split of a dataset and evaluated on the test-split. The only parameter tuned for the SVM is the L2 regularization strength, varied between 0.001 and 100. The overall test *accuracy* has been reported as a measure of performance.

## 5 Results and Discussion

The results of all our experiments are in compiled in Tables 3 and 4. All statistical tests reported are z-tests, where we compute the right-tailed p-value and call a result significantly different if  $p < 0.1$ .

**Performance of the distance metrics:** From Table 3 it can be observed that the word mover’s distance consistently performs better than the others experimented with in this paper. WMD calculates the total effort of aligning two sentences, which seems to capture more useful information compared to the hausdorff metric’s worst-case effort of alignment. As for the energy distance, it calculates pairwise potentials amongst words within and between sentences, and may suffer if there are

Method	Parameter	Value(s) Tested
EMAP	<i>n_neighbors</i>	40
	<i>embedding_dim</i>	50, 100, 300, 1000
	<i>min_dist</i>	1.0, 1.5, 2.0
	<i>spread</i>	1.0, 2.5
	<i>n_iters</i>	1000
	<i>distance</i>	wmd, hausdorff, energy
kNN	<i>k</i>	1
	<i>distance</i>	wmd, hausdorff, energy
dct	<i>components</i>	1 through 6
eigensent	<i>components</i>	1 through 3
	<i>time_lag</i>	1, 2, 3, [1,2], [1,2,3], [1,2,3,4]
pmeans	<i>powers</i>	1, [1,2], [1,2,3], [1,2,3,4,5,6]
s-bert	<i>model</i>	bert-base-nli-mean-tokens

Table 2: **Hyperparameter values tested.** For *EMAP*, *n\_neighbors* refers to the size of local neighborhood used for manifold approximation, *embedding\_dim* is the fixed dimensionality of the generated sentence-embeddings, *min\_dist* is the minimum distance apart that points are allowed to be in the low dimensional representation, *spread* determines the scale at which embedded points will be spread out, *n\_iters* is the number of iterations that the UMAP algorithm is allowed to run, and finally, *distance* is one of the metrics proposed in Section 2.1. For the spectral decomposition based algorithms, *dct* and *eigensent*, *components* represents the number of components to keep in the resulting decomposition, while *time\_lag* corresponds to the window-length in the dynamic mode decomposition process. For *pmeans*, *powers* represents the different powers which are used to generate the concatenated embeddings.

Distance	energydist		hausdorffdist		wmdist	
Method	knn	EMAP	knn	EMAP	knn	EMAP
amazon	<b>0.923*</b>	0.909	0.781	<b>0.844*</b>	0.918	<b>0.929*</b>
bbsport	0.941	<b>0.942</b>	0.925	<b>0.941</b>	0.972	<b>0.987</b>
classic	0.912	<b>0.921</b>	0.943	<b>0.953*</b>	0.961	<b>0.978*</b>
ohsumed	0.456	<b>0.505*</b>	0.491	<b>0.603*</b>	0.551	<b>0.630*</b>
r8	0.942	<b>0.962*</b>	<b>0.863*</b>	0.837	0.951	<b>0.973*</b>
twitter	0.731	<b>0.749</b>	0.736	<b>0.741</b>	0.712	<b>0.722</b>

Table 3: **Comparison versus kNN.** Results shown here compare the classification accuracies of *k*-nearest neighbour to our proposed approach for various distance metrics. For every distance, **bold** indicates better accuracy, while \* indicates that the winning accuracy was statistically significant with respect to the compared value (i.e., EMAP vs kNN for a given distance metric). It can be observed that our method almost always outperforms *k*-nearest neighbour-based classification.

Method	wmd-EMAP	dct	eigensent	wmovers	pmeans	doc2vec	s-bert
amazon	0.929	0.932	0.902∇	<b>0.943</b> ∧	0.938	0.912∇	0.923
bbsport	<b>0.986</b>	0.972	0.968	0.982	0.981	0.979	<b>0.986</b>
classic	<b>0.978</b>	0.964	0.947∇	0.971	0.960	0.965	0.966
ohsumed	0.630	0.594∇	0.574∇	<b>0.645</b> ∧	0.614∇	0.598∇	0.556∇
r8	<b>0.973</b>	0.967	0.958∇	0.972	0.969	0.949∇	0.954∇
twitter	0.722	0.644∇	0.669∇	<b>0.745</b>	0.636∇	0.673∇	0.673∇

Table 4: **Comparison versus competing methods.** We compare *EMAP* based on word mover’s distance to various state-of-the-art approaches. The best and second-best classification accuracies are highlighted in **bold** and *italics*. We perform statistical significance tests of our method (*wmd-EMAP*) against all other methods, for a given dataset, and denote the outcomes by ∇ when the compared method is worse and ∧ when our method is worse, while the absence of a symbol indicates insignificant differences. In terms of absolute accuracy, we observe that our method achieves state-of-the-art results in 2 out of 6 datasets.

shared commonly-occurring words in both the sentences. However, given that energy and hausdorff distances are reasonably fast to calculate and perform respectably well, they might be worth using in applications with a large number of long sentences.

**Comparison versus kNN:** *EMAP* almost always

outperforms *k*-nearest neighbours based classification, for all the tested distance metrics. The performance boost for WMD is between a relative percentage accuracy of 0.5% to 14%. This illustrates the efficiency of the proposed manifold-learning method.

Query Sentence	Best Match Sentence	Cosine Sim
I have spent thousands of dollar's On Meyers cookware everthing from KitchenAid Anolon Prestige Faberware & Circulan just to name a few Though Meyers does manufacture very high quality pots & pans and I would recommend them to anyone it's just sad that if you have any problem with them under warranty you have to go throught the chain of command that never gets you anywhere even if you want to speak with upper management about the rudeness of the customer service department Their customer service department employees are always very rude and snotty and they act like they are doing you a favor to even talk to you about their products	When I opened the box I noticed corrosion on the lid When I contacted Rival customer service via email they told me I had to purchase a new lid I called and spoke with a customer service representative and they told me that a lid was not covered under warranty When I explained that I just opened it and it was defective they told me to just return the product that there was nothing that they were going to do After being treated this way I will NOT be purchasing any more Rival products if they don't stand behind their product VERY VERY poor customer service	0.997
This movie will bring up your racial prejudices in ways that most movies just elude to It demonstrates how connected we all are as people and how seperated we are by only one thing our viewpoints The acting is superb and you get one cameo appearance after another which is a treat Of course the soundtrack is terrific The ending is intense to witness one situation after another coming to an unfortunate finish	I waited years for this movie to be released in the United States As far as I was concerned it wasn't about the acting as much as it was about the feeling the actors wanted to portray in which they profoundly accomplished I would recommend this movie to anyone who can reach that one step deeper into the minds of creativity and passion and appreciate the struggles of rising above and beyond the pain of broken dreams	0.998
We see a phrase a lot when we visit how to sites for writers World building By this we mean the setting the characters and everything else where our story will occur For me this often means maps memories and visits since I write about where I live But if you'd like to see exactly what world building means head down to your local library and grab SALEM'S LOT by Stephen King When Stephen King mania first gripped the English speaking world I missed it I saw the film of CARRIE and hated it Years later at a guard desk on a long shift scheduled so suddenly that I hadn't had a chance to visit the library I read what was in the desk instead THINNER If I were Stephen King I'd have put a pen name on that crap as well One of King's fans brought me around She recommended THE SHINING Of course I thought of that Kubrick/Nicholson travesty No no she said read the book It's much different Yes it is It's fantastic for its perceptiveness Next up PET SEMATARY which scared the crap out of me And that my friends is not easy ON WRITING I've gushed about that enough times The films STAND BY ME and THE APT PUPIL So in the end I appreciate King and forgive him for CARRIE and I think he's forgiven himself	in the possibility that Steve Berry could ever transcend his not so great debut The Amber Room Romanov Prophecy started in the right direction Third Secret was OK but I think he hit his *peak* right there	0.955

Table 5: **Examples of best-matching sentences.** From the *amazon* reviews dataset using *wmd-EMAP*.

### Comparison versus state-of-the-art methods:

Consulting Table 4, it seems that *wmovers*, *pmeans* and *s-bert* form the strongest baselines as compared to our method, *wmd-EMAP* (*EMAP* with word mover's distance). Considering the statistical significance of the differences in performance between *wmd-EMAP* and the others, it can be seen that it is almost always equivalent to or better than the other state-of-the-art approaches. In terms of absolute accuracy, it *wins* in 3 out of 6 evaluations, where it has the highest classification accuracy, and comes out second-best for the others. Compared

to it's closest competitor, the *word mover's embedding* algorithm, the performance of *wmd-EMAP* is found to be on-par (or slightly better, by 0.8% in the case of the *classic* dataset) to slightly worse (3% relative p.p., in case of the *twitter* dataset). Interestingly, both of the distance-based embedding approaches, *wmd-EMAP* and *wmovers*, are found to perform better than the siamese-BERT based approach, *s-bert*.

Thus, the overall conclusion from our empirical studies is that *EMAP* performs favourably as compared to various state-of-the-art approaches.



**Examples of similar sentences with EMAP:** We provide motivating examples of similar sentences from the *amazon* dataset, as deemed by our approach, in Table 5. As can be seen, our method performs quite well in matching complex sentences with varying topics and sentiments to their closest pairs. The first example pair has the theme of a customer who is unhappy about poor customer service in the context of cookware warranty, while the second one is about positive reviews of deeply-moving movies. The third example, about book reviews, is particularly interesting: in the first example, a reviewer is talking about how she disliked the first Stephen King work which she was exposed to, but subsequently liked all the next ones, while in the matched sentence the reviewer talks about a similar sentiment change towards the works of another author, Steve Berry. Thus in the last example, the similarity between sentences is the change of sentiment, from negative to positive, towards the works of books of particular authors.

## 6 Conclusions

In this work, we propose a novel mechanism to construct unsupervised sentence-embeddings by preserving properties of local neighbourhoods in the original space, as delineated by set-distance metrics. This method, which we term, *EMAP* or *Embeddings by Manifold Approximation and Projection* leverages a method from topological data analysis can be used as a framework with any distance metric that can discriminate between sets, three of which we test in this paper. Using both quantitative empirical studies, where we compare with state-of-the-art approaches, and qualitative probing, where we retrieve similar sentences based on our generated embeddings, we illustrate the efficiency of our proposed approach to be on-par or exceeding in-use methods. This work demonstrates the successful application of topological data analysis in sentence embedding creation, and we leave the design of better distance metrics and manifold approximation algorithms, particularly targeted towards NLP, for future research.

## References

Hanan Aldarmaki and Mona Diab. 2018. [Evaluation of unsupervised compositional representations](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2666–2677, Santa Fe, New Mexico, USA.

Nada Almarwani, Hanan Aldarmaki, and Mona Diab. 2019. [Efficient sentence embedding using discrete cosine transform](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3672–3678, Hong Kong, China.

Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. [A simple but tough-to-beat baseline for sentence embeddings](#). In *International Conference on Learning Representations*.

Mikhail J. Atallah. 1983. [A linear time algorithm for the hausdorff distance between convex polygons](#). Technical report, Department of Computer Science, Purdue University.

Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. [A neural probabilistic language model](#). *Journal of Machine Learning Research*, 3:1137–1155.

Karol Borsuk. 1948. On the imbedding of systems of compacta in simplicial complexes. In *Fundamenta Mathematicae*, volume 35, pages 217–234.

Gunnar Carlsson. 2009. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308.

Minmin Chen. 2017. [Efficient vector representation for documents through corruption](#). *5th International Conference on Learning Representations*.

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*, pages 670–680, Copenhagen, Denmark.

Corinna Cortes and Vladimir Vapnik. 1995. [Support-vector networks](#). *Machine Learning*, 20(3):273–297.

T. Cover and P. Hart. 1967. [Nearest neighbor pattern classification](#). *IEEE Transactions on Information Theory*, 13(1):21–27.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.

Wei Dong, Charikar Moses, and Kai Li. 2011. [Efficient k-nearest neighbor graph construction for generic similarity measures](#). In *Proceedings of the 20th International Conference on World Wide Web*, pages 577–586.

- M. . Dubuisson and A. K. Jain. 1994. [A modified hausdorff distance for object matching](#). In *Proceedings of 12th International Conference on Pattern Recognition*, volume 1, pages 566–568 vol.1.
- Geoffrey E Hinton and Sam T. Roweis. 2003. [Stochastic neighbor embedding](#). In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 857–864.
- Subhradeep Kayal and George Tsatsaronis. 2019. [EigenSent: Spectral sentence embeddings using higher-order dynamic mode decomposition](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4536–4546, Florence, Italy.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Skip-thought vectors](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3294–3302.
- Matt J. Kusner, Yu Sun, Nicholas I. Kolkin, and Kilian Q. Weinberger. 2015. [From word embeddings to document distances](#). In *Proceedings of the 32nd International Conference on Machine Learning - Volume 37, ICML’15*, pages 957–966.
- Quoc Le and Tomas Mikolov. 2014. [Distributed representations of sentences and documents](#). In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML’14*, pages II–1188–II–1196.
- Soledad Le Clainche and José M. Vega. 2017. [Higher order dynamic mode decomposition](#). *SIAM Journal on Applied Dynamical Systems*, 16(2):882–925.
- John A. Lee and Michel Verleysen. 2007. *Nonlinear Dimensionality Reduction*, 1st edition.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-SNE](#). *Journal of Machine Learning Research*, 9:2579–2605.
- L. McInnes, J. Healy, and J. Melville. 2018. [UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction](#). *ArXiv e-prints*.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. [Umap: Uniform manifold approximation and projection](#). *The Journal of Open Source Software*, 3(29):861.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13*, pages 3111–3119.
- H. Palangi, L. Deng, Y. Shen, J. Gao, X. He, J. Chen, X. Song, and R. Ward. 2016. [Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(4):694–707.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. [Scikit-learn: Machine learning in python](#). *J. Mach. Learn. Res.*, 12:2825–2830.
- 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.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- Ali Rahimi and Benjamin Recht. 2008. [Random features for large-scale kernel machines](#). In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1177–1184.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China.
- Lior Rokach and Oded Maimon. 2005. [Clustering methods](#). In *The Data Mining and Knowledge Discovery Handbook*, pages 321–352.
- Sam T. Roweis and Lawrence K. Saul. 2000. [Nonlinear dimensionality reduction by locally linear embedding](#). *Science*, 290:2323–2326.
- Andreas Rücklé, Steffen Eger, Maxime Peyrard, and Iryna Gurevych. 2018. [Concatenated p-mean word embeddings as universal cross-lingual sentence representations](#). *CoRR*, abs/1803.01400.
- Sidak Pal Singh, Andreas Hug, Aymeric Dieuleveut, and Martin Jaggi. 2019. [Context mover’s distance & barycenters: Optimal transport of contexts for building representations](#). In *Deep Generative Models for Highly Structured Data, ICLR 2019 Workshop, New Orleans, Louisiana, United States, May 6, 2019*.

- Gábor J. Székely and Maria L. Rizzo. 2013. [Energy statistics: A class of statistics based on distances](#). *Journal of Statistical Planning and Inference*, 143(8):1249 – 1272.
- George Tsatsaronis, Iraklis Varlamis, and Kjetil Nørvåg. 2012. [Semafor: Semantic document indexing using semantic forests](#). In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management, CIKM '12*, page 1692–1696, New York, NY, USA.
- Zihao Wang, Datong Zhou, Yong Zhang, Hao Wu, and Chenglong Bao. 2019. [Wasserstein-fisher-rao document distance](#). *CoRR*, abs/1904.10294.
- Lingfei Wu, Ian En-Hsu Yen, Kun Xu, Fangli Xu, Avinash Balakrishnan, Pin-Yu Chen, Pradeep Ravikumar, and Michael J. Witbrock. 2018. [Word mover’s embedding: From Word2Vec to document embedding](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4524–4534, Brussels, Belgium.
- Zhao Yan, Nan Duan, Junwei Bao, Peng Chen, Ming Zhou, Zhoujun Li, and Jianshe Zhou. 2016. [Doc-Chat: An information retrieval approach for chatbot engines using unstructured documents](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 516–525, Berlin, Germany.