

Correlations between Word Vector Sets

Vitalii Zhelezniak, April Shen, Daniel Busbridge,
Aleksandar Savkov, Nils Hammerla

{firstname.lastname}@babylonhealth.com

Abstract

Similarity measures based purely on word embeddings are comfortably competing with much more sophisticated deep learning and expert-engineered systems on unsupervised semantic textual similarity (STS) tasks. In contrast to commonly used geometric approaches, we treat a single word embedding as e.g. 300 observations from a scalar random variable. Using this paradigm, we first illustrate that similarities derived from elementary pooling operations and classic correlation coefficients yield excellent results on standard STS benchmarks, outperforming many recently proposed methods while being much faster and trivial to implement. Next, we demonstrate how to avoid pooling operations altogether and compare sets of word embeddings directly via correlation operators between reproducing kernel Hilbert spaces. Just like cosine similarity is used to compare individual word vectors, we introduce a novel application of the centered kernel alignment (CKA) as a natural generalisation of squared cosine similarity for sets of word vectors. Likewise, CKA is very easy to implement and enjoys very strong empirical results.

1 Introduction

Distributed representations of text have had a massive impact on the natural language processing (NLP), information retrieval (IR), and machine learning (ML) communities, thanks in part to their ability to capture rich notions of semantic similarity. While this work originally began with word embeddings (Bengio et al., 2003; Mikolov et al., 2013a; Pennington et al., 2014; Bojanowski et al., 2017; Joulin et al., 2017), there is now an ever-increasing number of representations for longer units of text based on simple aggregations of word vectors (Mitchell and Lapata, 2008; De Boom et al., 2016; Arora et al., 2017; Wieting

et al., 2016; Wieting and Gimpel, 2018; Zhelezniak et al., 2019b) as well as complex neural architectures (Le and Mikolov, 2014; Kiros et al., 2015; Hill et al., 2016; Conneau et al., 2017; Gan et al., 2017; Tang et al., 2017; Pagliardini et al., 2018; Zhelezniak et al., 2018; Subramanian et al., 2018; Cer et al., 2018; Devlin et al., 2018).

By contrast, relatively little effort has been directed towards understanding the similarity measures used to compare these textual embeddings, for which cosine similarity remains a convenient and widespread, yet somewhat arbitrary default, despite some emerging research into the alternatives (Camacho-Collados et al., 2015; De Boom et al., 2015; Santus et al., 2018; Zhelezniak et al., 2019b,a). Part of the appeal of cosine similarity perhaps lies in the simple geometric interpretation behind it. However, as embeddings are ultimately just arrays of numbers, we are free to take alternative viewpoints other than the geometric ones, if they lead to illuminating insights or strong-performing methods.

Following Zhelezniak et al. (2019a), we treat a word embedding not as a geometric vector but as a statistical sample (of e.g. 300 observations) from a scalar random variable, and indeed find insights that are both intriguing and noteworthy. We first illustrate that similarities derived from elementary pooling operations and classic univariate correlation coefficients yield excellent results on standard semantic textual similarity (STS) benchmarks, outperforming many recently proposed methods while being much faster and simpler to implement. This empirically validates the advantages of the statistical perspective on word embeddings over the geometric interpretations. In the process, we provide more evidence that departures from normality, and in particular the presence of outliers, can have severe negative effects on the performance of some correlation coeffi-

cients. We show how to overcome these complications, by selecting an outlier-removing pooling operation such as max-pooling, applying a more robust correlation coefficient such as Spearman’s ρ , or simply clipping (winsorizing) the word vectors.

Next, we demonstrate how to avoid pooling operations completely and compare sets of word embeddings directly via correlation operators between reproducing kernel Hilbert spaces (RKHS). We introduce a novel application of the kernel alignment (KA) and the centered kernel alignment (CKA) as a natural generalisation of the squared cosine similarity and Pearson correlation for the sets of word embeddings. These multivariate correlation coefficients are very easy to implement and also enjoy very strong empirical results.

2 Related Work

Several lines of research seek to combine the strength of pretrained word embeddings and the elegance of set- or bag-of-words (BoW) representations. Any method that determines semantic similarity between sentences by comparing the corresponding sets of word embeddings is directly related to our work.

Perhaps the most obvious such approaches are based on elementary pooling operations such as average-, max- and min-pooling (Mitchell and Lapata, 2008; De Boom et al., 2015, 2016). While seemingly over-simplistic, numerous studies have confirmed their impressive performance on the downstream tasks (Arora et al., 2017; Wieting et al., 2016; Wieting and Gimpel, 2018; Zhelezniak et al., 2019b)

One step further, Zhao and Mao (2017); Zhelezniak et al. (2019b) introduce fuzzy bags-of-words (FBoW) where degrees of membership in a fuzzy set are given by the similarities between word embeddings. Zhelezniak et al. (2019b) show a close connection between FBoW and max-pooled word vectors.

Some approaches do not seek to build an explicit representation and instead focus directly on designing a similarity function between sets. Word Mover’s Distance (WMD) (Kusner et al., 2015) is an instance of the Earth Mover’s Distance (EMD) computed between normalised BoW, with the cost matrix given by Euclidean distances between word embeddings. In the soft cardinality framework of (Jimenez et al., 2010, 2015), the

contribution of a word to the cardinality of a set depends on its similarities to other words in the same set. Such sets are then compared using an appropriately defined Jaccard index or related measures. DynaMax (Zhelezniak et al., 2019b) uses universe-constrained fuzzy sets designed explicitly for similarity computations.

Approaches that see word embeddings as statistical objects are very closely related to our work. Virtually all of them treat word embeddings as observations from some D -variate parametric family, where D is the embedding dimension. Arora et al. (2016, 2017) introduce a latent discourse model and show the maximum likelihood estimate (MLE) for the discourse vector to be the weighted average of word embeddings in a sentence, where the weights are given by smooth inverse frequencies (SIF). Nikolentzos et al. (2017); Torki (2018) treat sets of word embeddings as observations from D -variate Gaussians, and compare such sets with cosine similarity between the parameters (means and covariances) estimated by maximum likelihood. Vargas et al. (2019) measure semantic similarity through penalised likelihood ratio between the joint and factorised models and explore Gaussian and von Mises–Fisher likelihoods.

Cosine similarity between covariances is an instance of the RV coefficient and its uncentered version was applied in the context of word embeddings before (Botev et al., 2017). We arrive at a similar coefficient (but with different centering) as a special case of CKA, which in the general case makes no parametric assumptions about distributions whatsoever. In particular our version is suitable for comparing sets containing just one word vector, whereas the method of Nikolentzos et al. (2017); Torki (2018) requires at least two vectors in each set. Very recently, Kornblith et al. (2019) used CKA to compare representations between layers of the same or different neural networks. This is again an instance of treating such representations as observations from a D -variate distribution, where D is the dimension of the hidden layer in question. Our use of CKA is completely different from theirs.

Unlike all of the above approaches, (Zhelezniak et al., 2019a) see each word embedding itself as D (e.g. 300) observations from some scalar random variable. They cast semantic similarity as correlations between these random variables and study

their properties using simple tools from univariate statistics. While they consider correlations between individual word vectors and averaged word vectors, they do not formally explore correlations between word vector sets. We review their framework in Section 3 and then proceed to formalise and generalise it to the case of sets of word embeddings.

3 Background: Correlation Coefficients and Semantic Similarity

Suppose we have a word embeddings matrix $\mathbf{W} \in \mathbb{R}^{N \times D}$, where N is the number of words in the vocabulary and D is the embedding dimension (usually 300). In other words, each row $\mathbf{w}^{(i)}$ of \mathbf{W} is a D -dimensional word vector. When applying statistical analysis to these vectors, one might choose to treat each $\mathbf{w}^{(i)}$ as an observation from some D -variate distribution $P_D(E_1, \dots, E_D)$ and model it with a Gaussian or a Gaussian Mixture. While such analysis helps in studying the overall geometry of the embedding space (how dimensions correlate and how embeddings cluster), P_D is not directly useful for semantic similarity between individual words.

For the latter, Zhelezniak et al. (2019a) proposed to look at the transpose \mathbf{W}^T and the corresponding distribution $P(W_1, W_2, \dots, W_N)$. Under this perspective, each word vector $\mathbf{w}^{(i)}$ is now a sample of D (e.g. 300) observations from a scalar random variable W_i . Luckily, in applications we are usually not interested in the full joint distribution but only in the similarity between two words, i.e. the bivariate marginal $P(W_i, W_j)$. In practice, we make inferences about this marginal from the paired sample $(\mathbf{w}^{(i)}, \mathbf{w}^{(j)})$ through visualisations (histograms, Q-Q plots, scatter plots, etc.) as well as various statistics.

Zhelezniak et al. (2019a) found that for all common models (GloVe, fastText, word2vec) the means across word embeddings are tightly concentrated around zero (relative to their dimensions), thus making the widely used cosine similarity practically equivalent to Pearson correlation. However, while word2vec vectors seem mostly normal, GloVe and fastText vectors are highly non-normal, likely due to the presence of heavy univariate and bivariate outliers (as suggested by visualisations mentioned earlier). Quantitatively, the majority of GloVe and fastText vectors fail the Shapiro-Wilk normality test at sig-

nificance level 0.05. Therefore, while Pearson’s r (and thus cosine similarity) may be acceptable for word2vec, it is preferable to resort to more robust non-parametric correlation coefficients such as Spearman’s ρ or Kendall’s τ as a similarity measure between GloVe and fastText vectors.

Finally, very similar conclusions were shown to hold for sentence representations obtained by word vector averaging, also referred to as mean-pooling. In particular, averaged fastText vectors compared with rank correlation coefficients already show impressive results on standard STS tasks, rivaling much more sophisticated systems.

4 Correlations between Word Vector Sets

We are interested in applying the statistical framework from Section 3 to measure the semantic similarity between two sentences s_1 and s_2 given by the sets (or bags) S_1 and S_2 of word embeddings respectively. To formalise this new setup, we may see each set of word embeddings $S = \{\mathbf{w}_{(1)}, \mathbf{w}_{(2)}, \dots, \mathbf{w}_{(k)}\}$ as a sample (of e.g. 300 observations) from some theoretical set of scalar random variables $R = \{W_1, W_2, \dots, W_k\}$. In light of the above, our task then lies in finding correlation coefficients $\text{corr}(R_1, R_2)$ between R_1 and R_2 and their empirical estimates $\widehat{\text{corr}}(S_1, S_2)$ obtained from the paired sample S_1, S_2 , hoping that such coefficients will serve as a good proxy for semantic similarity. Recall that for single-word sets $R_1 = \{W_i\}, R_2 = \{W_j\}$ the task simply reduces to computing a univariate correlation between word vectors $\mathbf{w}^{(i)}$ and $\mathbf{w}^{(j)}$, where the choice of the coefficient (Pearson’s r , Spearman’s ρ , etc.) is made based on the statistics exhibited by the word embeddings matrix. While generalising this to sets of more than one variable is not particularly hard, there are several ways to do so, each with its own advantages and downsides. In the present work, we group these approaches into two broad families: pooling-based and pooling-free correlation coefficients.

4.1 Correlations between Pooled Variables

Pooling-based approaches first reduce a set of random variables to a single scalar random variable $W_{\text{pool}} = f_{\text{pool}}(W_1, W_2, \dots, W_k)$ and then apply univariate correlation coefficients between the pooled variables. In practice this would correspond to pooling word embeddings $\mathbf{w}_{(1)}, \mathbf{w}_{(2)}, \dots, \mathbf{w}_{(k)}$ (along $i = 1:k$) into one




Figure 1: Normalised histograms of the mean distribution for sentence vectors generated by mean-, max- and min-pooling. Sentences were taken from the entire STS dataset (Agirre et al., 2012, 2013, 2014, 2015, 2016; Cer et al., 2017), and we utilise three commonly-used word embedding models: GloVe (Pennington et al., 2014), fastText (Bojanowski et al., 2017), and word2vec (Mikolov et al., 2013b,c).

Algorithm 1 MaxPool-Spearman

Require: Word embeddings for the first sentence

$$\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(k)} \in \mathbb{R}^{1 \times d}$$

Require: Word embeddings for the second sentence

$$\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(l)} \in \mathbb{R}^{1 \times d}$$

Ensure: Similarity score MS

Max-pooling performed element-wise

$$\mathbf{x} \leftarrow \text{MAX_POOL}(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(k)})$$

$$\mathbf{y} \leftarrow \text{MAX_POOL}(\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(l)})$$

$$MS \leftarrow \text{SPEARMANCORRELATION}(\mathbf{x}, \mathbf{y})$$

fixed vector \mathbf{w}_{pool} , followed by computing univariate sample correlations. Certainly, these approaches are empirically attractive: not only are they very simple computationally (e.g. see Algorithm 1) but they also keep us in the realm of univariate statistics, where we have an entire arsenal of effective tools for making inferences about W_{pool} .

Unfortunately, it is not always clear *a priori* what should dictate our choice of the pooling function (though, as we will see shortly, for certain functions some statistical justifications do exist). By far the most common pooling operations for word embedding found in the literature are mean-, max- and min-pooling. It is also very common, with some exceptions, to treat these various pooled representation in a completely identical fashion,

e.g. by comparing them all with cosine similarity. Intuitively, however, we suggest that the statistics of W_{pool} must heavily depend on the pooling function f_{pool} and thus each such pooled random variable should be studied in its own right. To illustrate this point, we would like to reveal the very different nature of mean- and max- and min-pooled sentence vectors though a practical example.

4.2 Statistics of the Pooled Representations: A Practical Analysis

Let us begin by examining sentence vectors obtained through mean-pooling. Recall that for common word embedding models, the mean across 300 dimensions of a single word embedding $\mathbf{w}_{(i)}$ happens to be close to zero (relative to the dimensions). By the linearity of expectation, we have that $\mathbb{E}[W_{\text{mean}}] = \mathbb{E}\left[\sum_{i=1}^k W_i\right] = \sum_{i=1}^k \mathbb{E}[W_i]$, and so the mean across \mathbf{w}_{mean} will also be close to zero at least for small k . In practice, this seems to hold even for moderate k in naturally occurring sentences, as seen in Figure 1. Based on this, we expect Pearson correlation and cosine similarity to have almost identical performance on the downstream tasks, which is confirmed in Figure 2.

On the other hand, intuition tells us that the means of the max-pooled vectors will be shifted




Figure 2: Bar plots of Pearson correlation on STS tasks between human scores and the following set-based similarity metrics: Cosine similarity (COS), Pearson’s r (PRS), Winsorized Pearson’s r (WPRS), Spearman’s ρ (SPR), and Kendall’s τ (KND). Plots generated for three pooling methods and the following word embedding models: GloVe (Pennington et al., 2014), fastText (Bojanowski et al., 2017), and word2vec (Mikolov et al., 2013b,c).

to the right because of the max operation, which we see in Figure 1. In this case, cosine similarity and Pearson correlation will yield different results and, in fact, Pearson’s r considerably outperforms cosine on the downstream tasks (Figure 2). This in turn empirically adds weight to the statistical interpretation (correlation) over its geometrical counterpart (angle between vectors).

Recall also that unlike word2vec, GloVe and fastText vectors feature heavy univariate outliers, and the same can be expected to hold for the pooled representations; an example is shown in Figure 3. In case of mean-pooled vectors, this particular departure from normality can be successfully detected by the Shapiro-Wilk normality test, informing the appropriate choice of the correlation coefficient (Pearson’s r or robust rank correlation). By contrast, such procedure cannot be readily applied to max-pooled and min-pooled vectors as by construction they exhibit additional departures from normality, such as positive and negative skew respectively. It is always a good idea to consult visualisations for such vectors, such as the ones in Figure 3. Interestingly though, we do observe the some noteworthy regularities, which we describe further in Section 5.

The above example is meant to illustrate that even the simplest pooled random variables show

strikingly different statistics depending on the aggregation. While the abundance of various pooling operations may be intimidating, the resulting vectors are always subject to the many tools of univariate statistics. As we hope to have shown, even crude analysis can shed light on the nature of these textual representations, which in turn has notable practical implications, as we will see in Section 5.

4.3 Correlations between Random Vectors

Exactly as before, suppose we have two sentences $S_1 = \{\mathbf{x}_{(1)}, \mathbf{x}_{(2)}, \dots, \mathbf{x}_{(k)}\}$ and $S_2 = \{\mathbf{y}_{(1)}, \mathbf{y}_{(2)}, \dots, \mathbf{y}_{(l)}\}$ and the corresponding random vectors $\mathbf{X} = (X_1, X_2, \dots, X_k)$ and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_l)$. At this point it is important to emphasise again that we relate each word vector \mathbf{x}_i to a random variable X_i and treat the dimensions of \mathbf{x}_i as D observations from that variable, and similarly for \mathbf{y}_i and Y_i . In contrast with the pooling-based approaches, our task here is to find a suitable correlation coefficient directly between the random vectors \mathbf{X} and \mathbf{Y} . We begin by recalling the expression for the basic univariate Pearson’s r :

$$r_{XY} = \frac{\mathbb{E}_{X,Y} [(X - \mu_X)(Y - \mu_Y)]}{\sigma_X \sigma_Y}, \quad (1)$$




Figure 3: Histograms for word embeddings of the word “cats” and pooled representations of the embeddings for the words in the sentence “I like cats because they are very cute animals”.

where

$$\mu_X = \mathbb{E}[X], \quad \sigma_X = \sqrt{\mathbb{E}[X^2] - \mu_X^2},$$

and similarly for μ_Y and σ_Y . The covariance term $\text{cov}(X, Y)$ in the numerator is readily generalised to random vectors by the following cross-covariance operator between reproducing kernel Hilbert spaces (RKHS) \mathcal{F} and \mathcal{G}

$$\mathbf{C}_{\mathbf{X}\mathbf{Y}} = \mathbb{E}_{\mathbf{X}, \mathbf{Y}} [(\phi(\mathbf{X}) - \mu_{\mathbf{X}}) \otimes (\psi(\mathbf{Y}) - \mu_{\mathbf{Y}})], \quad (2)$$

where \otimes denotes the tensor product and $\mu_{\mathbf{X}} = \mathbb{E}_{\mathbf{X}}[\phi(\mathbf{X})]$, $\mu_{\mathbf{Y}} = \mathbb{E}_{\mathbf{Y}}[\psi(\mathbf{Y})]$. Here ϕ and ψ are the feature maps such that $\langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathcal{F}} = K(\mathbf{x}, \mathbf{x}')$ and $\langle \psi(\mathbf{y}), \psi(\mathbf{y}') \rangle_{\mathcal{G}} = L(\mathbf{y}, \mathbf{y}')$, where K and L are the kernels associated with RKHS \mathcal{F} and \mathcal{G} respectively. Note that if ϕ and ψ are the identity maps, the cross-covariance operator (2) simply becomes the cross-covariance matrix

$$\mathbf{C}_{\mathbf{X}\mathbf{Y}} = \mathbb{E}_{\mathbf{X}, \mathbf{Y}} [(\mathbf{X} - \mu_{\mathbf{X}})(\mathbf{Y} - \mu_{\mathbf{Y}})^T].$$

Gretton et al. (2005a) define the Hilbert-Schmidt independence criterion (HSIC) to be the squared Hilbert-Schmidt norm $\|\mathbf{C}_{\mathbf{X}\mathbf{Y}}\|_{\text{HS}}^2$ of (2) and derive an expression for it in terms of kernels K and L

$$\begin{aligned} \text{HSIC}(\mathbf{X}, \mathbf{Y}, K, L) = & \mathbb{E}_{\mathbf{X}, \mathbf{X}', \mathbf{Y}, \mathbf{Y}'} [K(\mathbf{X}, \mathbf{X}')L(\mathbf{Y}, \mathbf{Y}')] \\ & + \mathbb{E}_{\mathbf{X}, \mathbf{X}'} [K(\mathbf{X}, \mathbf{X}')] \mathbb{E}_{\mathbf{Y}, \mathbf{Y}'} [L(\mathbf{Y}, \mathbf{Y}')] \\ & - 2\mathbb{E}_{\mathbf{X}, \mathbf{Y}} [\mathbb{E}_{\mathbf{X}'} [K(\mathbf{X}, \mathbf{X}')] \mathbb{E}_{\mathbf{Y}'} [L(\mathbf{Y}, \mathbf{Y}')]]. \end{aligned}$$

They also show the empirical estimate of it to be

$$\text{HSIC}(\mathbf{K}, \mathbf{L}) = (D - 1)^{-2} \text{Tr}(\mathbf{K}\mathbf{H}\mathbf{L}\mathbf{H}), \quad (3)$$

where $\mathbf{H} = \mathbf{I} - \frac{1}{D}\mathbf{1}\mathbf{1}^T$ is the centering matrix and $\mathbf{K} = K(\mathbf{X}^{(i)}, \mathbf{X}^{(j)})$, $\mathbf{L} = L(\mathbf{Y}^{(i)}, \mathbf{Y}^{(j)})$, $i, j = 1:D$ are the kernel (Gram) matrices of observations. Crucially, the kernel evaluations for \mathbf{K} take place between $\mathbf{X}^{(i)} = (\mathbf{x}_{(1)}^i, \mathbf{x}_{(2)}^i, \dots, \mathbf{x}_{(k)}^i)$ and $\mathbf{X}^{(j)} = (\mathbf{x}_{(1)}^j, \mathbf{x}_{(2)}^j, \dots, \mathbf{x}_{(k)}^j)$ and not between the individual word embeddings $\mathbf{x}_{(i)}$ and $\mathbf{x}_{(j)}$, and similarly for \mathbf{L} . Thus, both \mathbf{K} and \mathbf{L} are square matrices of dimension $D \times D$. Indeed, for (3) to make sense, the dimensions of \mathbf{K} and \mathbf{L} must match. The matching dimension in our case is the word embedding dimension D , while the number of words k and l in the sentences may vary. This is in line with our formalism, which models word vectors as random variables and their dimensions as observations.

Finally, the Centered Kernel Alignment (CKA) (Cortes et al., 2012) is simply defined as

$$\text{CKA}(\mathbf{K}, \mathbf{L}) = \frac{\text{HSIC}(\mathbf{K}, \mathbf{L})}{\sqrt{\text{HSIC}(\mathbf{K}, \mathbf{K})}\sqrt{\text{HSIC}(\mathbf{L}, \mathbf{L})}}. \quad (4)$$

We see now that CKA not only generalises the squared Pearson correlation to the multivariate case, it also allows it to operate in high-dimensional feature spaces, as commonly done in the kernel literature. The reason this is useful is that under certain conditions (when K and L are

Approach	STS	12	13	14	15	16
<i>Deep Learning Approaches</i>						
ELMo (BoW)		55	53	63	68	60
Skip-Thought		41	29	40	46	52
InferSent		61	56	68	71	71
USE (DAN)		59	59	68	72	70
USE (Transformer)		61	64	71	74	74
STN (multitask)		60.6	54.7 [†]	65.8	74.2	66.4
BERT-Base		46.9	52.8	57.2	63.5	64.5
BERT-Large		42.6	47.4	49.3	55.6	60.4
<i>Set-based Approaches</i>						
WMD		53.0	45.9	57.2	66.9	63.1
SoftCardinality		55.9	50.5	59.0	66.2	65.1
DynaMax		60.9	60.3	69.5	76.7	74.6
SIF+PCA		58.6	67.3	70.5	73.5	71.7
MeanPool+COS		58.3	57.9	64.9	67.6	64.3
MaxPool+COS		57.7	53.5	67.2	69.5	68.5
MeanPool+SPR		60.2	65.1	70.1	74.4	73.0
MaxPool+SPR [✕]		61.0	62.9	70.9	75.9	75.8
CKA Linear [✕]		59.8	62.1	69.5	74.6	70.3
CKA Gaussian [✕]		60.5	63.8	71.6	76.3	73.7
CKA dCor [✕]		61.0	63.2	71.5	75.6	72.4

Table 1: Mean Pearson correlation on STS tasks for Deep Learning and Set-based methods using fastText. Methods proposed in this work are denoted with ✕. Values in bold indicate best results per task. Previous results are taken from Perone et al. (2018), Subramanian et al. (2018) and Zhelezniak et al. (2019b,a). [†] indicates the only STS13 result (to our knowledge) that includes the SMT subtask.

characteristic kernels), HSIC can detect any existing dependence with high probability, as the sample size increases (Gretton et al., 2005b). One can also consider the Uncentered Kernel Alignment (or simply KA) (Cristianini et al., 2002), which can then be seen as a similar generalisation but for the univariate cosine similarity. To the best of our knowledge, KA and CKA in general have never been applied before to measure semantic similarity between sets of word embeddings; therefore this work seeks to introduce them as standard definitions for squared Pearson’s r and cosine similarity for such sets.

5 Experiments

We now empirically demonstrate the power of the methods and statistical analysis presented in Section 4, through a set of evaluations on the Se-

mantic Textual Similarity (STS) tasks series 2012-2016 (Agirre et al., 2012, 2013, 2014, 2015, 2016; Cer et al., 2017). For methods involving pre-trained word embeddings, we use fastText (Bojanowski et al., 2017) trained on Common Crawl (600B tokens), as previous evaluations have indicated that fastText vectors have uniformly the best performance on these tasks out of commonly-used pretrained unsupervised word vectors (Conneau et al., 2017; Perone et al., 2018; Zhelezniak et al., 2019a,b). We provide experiments and significance analysis for additional word vector in the Appendix. The success metric for the STS tasks is the Pearson correlation between the sentence similarity scores provided by human annotators and the scores generated by a candidate algorithm. Note that the dataset for the STS13 SMT subtask is no longer publicly available, so the mean Pearson correlation for STS13 reported in our experiments has been re-calculated accordingly. The code for our experiments builds on the SentEval toolkit (Conneau and Kiela, 2018) and is available on GitHub¹.

We first conduct a set of experiments to validate the observations of Sections 4.1 and 4.2 regarding the performance of cosine similarity and various univariate correlation coefficients when applied to pooled word vectors. These results are depicted in Figure 2, for which we can make the following observations.

First, max and min-pooled vectors consistently outperform mean-pooled vectors when all three representations are compared with Pearson correlation. We hypothesise that this is in part because max and min-pooling remove the outliers (to which Pearson’s r is very sensitive) from at least one tail of the distribution whereas mean-pooled vectors have outliers in both tails. This outlier-removing property, however, cannot be taken as a sole explanation behind excellent performance of max-pooled vectors, as max-pooling still tends to outperform mean-pooling when both are compared with correlations that are robust to outliers, as well as on word vectors that have very few outliers to begin with (e.g. word2vec).

In addition, the strong performance of rank correlation coefficients (Spearman’s ρ and Kendall’s τ) comes solely from their robustness to outliers, as clipping (winsorizing) the top and

¹<https://github.com/Babylonpartners/corrsim>

bottom 5% of the values and then proceeding with Pearson’s r closes the gap almost completely. Consistently, on vectors with few outliers (word2vec), Pearson’s r achieves the same performance as rank correlations even without winsorization. However, unlike outliers, positive (negative) skew of max- (min-) pooled vectors does not seem to hurt Pearson’s r on STS tasks.

Next, we conduct evaluations of the methods proposed in this work alongside other deep learning and set-based similarity measures for STS from the literature. The methods we compare are as follows:

- *Deep representation approaches:* BoW with ELMo embeddings (Peters et al., 2018), Skip-Thought (Kiros et al., 2015), InferSent (Conneau et al., 2017), Universal Sentence Encoder both DAN and Transformer (Cer et al., 2018), STN multitask embeddings (Subramanian et al., 2018), and BERT 12- and 24-layer models (Devlin et al., 2018).
- *Set-based similarity measures:* Word Mover’s Distance (WMD) (Kusner et al., 2015), soft-cardinality with Jaccard coefficient (Jimenez et al., 2012), DynaMax with Jaccard (Zhelezniak et al., 2019b), mean- and max-pooled word vectors with cosine similarity (COS), and mean-pooled word vectors with Spearman correlation (SPR) (Zhelezniak et al., 2019a).
- *Proposed set-based approaches:* max-pooled word vectors with Spearman correlation, CKA with linear kernel (also known as RV-coefficient), CKA with Gaussian kernel (median estimation for σ^2), and CKA with distance kernel (distance correlation).

Note that for BERT we evaluated all pooling strategies available in bert-as-service (Xiao, 2018) applied to either the last or second-to-last layers and report results for the best-performing combination, which was mean-pooling on the last layer for both model sizes. Our results are presented in Table 1. We can clearly see that deep learning-based methods do not shine on STS tasks, while simple compositions of word vectors can perform extremely well, especially when an appropriate correlation coefficient is used as the similarity measure. Indeed, the performance of max-pooled vectors with Spearman correlation approaches or

Approach	Time complexity
MaxPool+SPR	$O(nd + d \log d)$
CKA	$O(nd^2 + d^2)$
DynaMax	$O(n^2 d)$
SoftCard	$O(n^2 d)$
WMD	$O(n^3 \log n \cdot d)$
WMD (relaxed)	$O(n^2 d)$

Table 2: Computational complexity of some of the set-based STS methods discussed in this paper. Here n is the sentence length and d is the dimensionality of the word embeddings.

exceeds that of more expensive or offline methods like that of Arora et al. (2017), which performs PCA computations on the entire test set. Additionally, while the multivariate correlation methods such as CKA are more computationally expensive than pooling-based approaches (see Table 2), they can provide performance boost on some tasks, making the cost worth it depending on the application. Finally, we conducted an exploratory error analysis and found that many errors are due to the well-known inherent weaknesses of word embeddings. For example, the proposed approaches heavily overestimate similarity when two sentences contain antonyms or when one sentence is the negation of the other. We illustrate these and other cases in the Appendix.

6 Conclusion

In this work we investigate the application of statistical correlation coefficients to sets of word vectors as a method for computing semantic textual similarity (STS). This can be done either by pooling these word vectors and computing univariate correlations between the resulting representations, or by applying multivariate correlation coefficients to the sets of vectors directly.

We provide further empirical evidence that outliers in word vector distributions disrupt performance of set-based similarity metrics as previously shown (Zhelezniak et al., 2019a). We also show working methods for solving or avoiding the issue through vector pooling operations, robust correlations or winsorization. In addition, we found that pooling operations in conjunction with univariate correlation coefficients yield one of the strongest results on downstream STS tasks, while being computationally much more efficient than competing set-based methods. Our findings are

supported by a combination of statistical analysis, practical examples and visualisations, and empirical evaluation on standard benchmark datasets.

Both proposed families of approaches serve as strong baselines for future research into STS, as well as useful algorithms for the practitioner, being efficient and simple to implement.

We believe our findings speak to the efficacy of the statistical perspective on word embeddings, which we hope will encourage others to explore further implications of not only this particular framework, but also completely novel interpretations of textual representations.

Acknowledgements

We would like to thank the three anonymous reviewers for their useful feedback and suggestions.

References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Inigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, German Rigau, Larraitz Uria, and Janyce Wiebe. 2015. **Semeval-2015 task 2: Semantic textual similarity, english, spanish and pilot on interpretability**. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 252–263. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. **Semeval-2014 task 10: Multilingual semantic textual similarity**. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 81–91. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2016. **Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation**. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. **Semeval-2012 task 6: A pilot on semantic textual similarity**. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. ***sem 2013 shared task: Semantic textual similarity**. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 32–43. Association for Computational Linguistics.
- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. **A latent variable model approach to pmi-based word embeddings**. *Transactions of the Association for Computational Linguistics*, 4:385–399.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. **A Simple but Tough-to-Beat Baseline for Sentence Embeddings**. *International Conference on Learning Representations*.
- Yoshua Bengio, Rjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. **A neural probabilistic language model**. *Journal of Machine Learning Research*, 3:1137–1155.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. **Enriching word vectors with subword information**. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Viktor Botev, Kaloyan Marinov, and Florian Schäfer. 2017. **Word importance-based similarity of documents metric (wisdm): Fast and scalable document similarity metric for analysis of scientific documents**. In *Proceedings of the 6th International Workshop on Mining Scientific Publications, WOSP 2017*, pages 17–23, New York, NY, USA. ACM.
- José Camacho-Collados, Mohammad Taher Pilehvar, and Roberto Navigli. 2015. **Nasari: a novel approach to a semantically-aware representation of items**. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 567–577. Association for Computational Linguistics.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. **Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation**. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14. Association for Computational Linguistics.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. **Universal sentence encoder for english**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174. Association for Computational Linguistics.

- Alexis Conneau and Douwe Kiela. 2018. [Senteval: An evaluation toolkit for universal sentence representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*. European Language Resource Association.
- 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. Association for Computational Linguistics.
- Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. 2012. Algorithms for learning kernels based on centered alignment. *Journal of Machine Learning Research*, 13(Mar):795–828.
- Nello Cristianini, John Shawe-Taylor, Andre Elisseeff, and Jaz S Kandola. 2002. On kernel-target alignment. In *Advances in neural information processing systems*, pages 367–373.
- C. De Boom, S. Van Canneyt, S. Bohez, T. Demeester, and B. Dhoedt. 2015. [Learning semantic similarity for very short texts](#). In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, pages 1229–1234.
- Cedric De Boom, Steven Van Canneyt, Thomas Demeester, and Bart Dhoedt. 2016. [Representation learning for very short texts using weighted word embedding aggregation](#). *Pattern Recogn. Lett.*, 80(C):150–156.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Zhe Gan, Yunchen Pu, Ricardo Henao, Chunyuan Li, Xiaodong He, and Lawrence Carin. 2017. [Learning generic sentence representations using convolutional neural networks](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2390–2400. Association for Computational Linguistics.
- Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. 2005a. Measuring statistical dependence with hilbert-schmidt norms. In *International conference on algorithmic learning theory*, pages 63–77. Springer.
- Arthur Gretton, Ralf Herbrich, Alexander Smola, Olivier Bousquet, and Bernhard Schölkopf. 2005b. Kernel methods for measuring independence. *Journal of Machine Learning Research*, 6(Dec):2075–2129.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. [Learning distributed representations of sentences from unlabelled data](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1367–1377. Association for Computational Linguistics.
- Sergio Jimenez, Claudia Becerra, and Alexander Gelbukh. 2012. [Soft cardinality: A parameterized similarity function for text comparison](#). In *Proceedings of the First Joint Conference on Lexical and Computational Semantics - Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation, SemEval '12*, pages 449–453, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Sergio Jimenez, Fabio Gonzalez, and Alexander Gelbukh. 2010. Text comparison using soft cardinality. In *String Processing and Information Retrieval*, pages 297–302, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Sergio Jimenez, Fabio A. Gonzalez, and Alexander Gelbukh. 2015. [Soft cardinality in semantic text processing: Experience of the SemEval international competitions](#). *Polibits*, 51:63–72.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431. Association for Computational Linguistics.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Skip-Thought Vectors](#). In *Advances in Neural Information Processing Systems*, pages 3294–3302.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. Similarity of neural network representations revisited. In *ICML*.
- 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 International Conference on Machine Learning*, volume 37 of *ICML'15*, pages 957–966. JMLR.org.
- Quoc Le and Tomas Mikolov. 2014. [Distributed Representations of Sentences and Documents](#). In *International Conference on Machine Learning*, pages 1188–1196.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient Estimation of Word Representations in Vector Space](#). *arXiv preprint arXiv:1301.3781*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems*, pages 3111–3119.

- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013c. **Linguistic Regularities in Continuous Space Word Representations**. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 746–751.
- Jeff Mitchell and Mirella Lapata. 2008. **Vector-based models of semantic composition**. In *Proceedings of ACL-08: HLT*, pages 236–244. Association for Computational Linguistics.
- Giannis Nikolentzos, Polykarpos Meladianos, Francois Rousseau, Yannis Stavrakas, and Michalis Vazirgiannis. 2017. **Multivariate Gaussian document representation from word embeddings for text categorization**. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 450–455, Valencia, Spain. Association for Computational Linguistics.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. **Unsupervised learning of sentence embeddings using compositional n-gram features**. 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 528–540. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Christian S Perone, Roberto Silveira, and Thomas S Paula. 2018. **Evaluation of sentence embeddings in downstream and linguistic probing tasks**. *arXiv preprint arXiv:1806.06259*.
- 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. Association for Computational Linguistics.
- Enrico Santus, Hongmin Wang, Emmanuele Chersoni, and Yue Zhang. 2018. **A rank-based similarity metric for word embeddings**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 552–557. Association for Computational Linguistics.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. **Learning general purpose distributed sentence representations via large scale multi-task learning**. In *International Conference on Learning Representations*.
- Shuai Tang, Hailin Jin, Chen Fang, Zhaowen Wang, and Virginia R. de Sa. 2017. **Exploring asymmetric encoder-decoder structure for context-based sentence representation learning**. *CoRR*, abs/1710.10380.
- Marwan Torki. 2018. **A document descriptor using covariance of word vectors**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 527–532, Melbourne, Australia. Association for Computational Linguistics.
- Francisco Vargas, Kamen Brestnichki, and Nils Hammerla. 2019. **Model comparison for semantic grouping**. In *ICML*.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2016. **Towards Universal Paraphrastic Sentence Embeddings**. In *International Conference on Learning Representations*.
- John Wieting and Kevin Gimpel. 2018. **Paranmt-50m: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462. Association for Computational Linguistics.
- Han Xiao. 2018. **bert-as-service**. <https://github.com/hanxiao/bert-as-service>.
- Rui Zhao and Kezhi Mao. 2017. **Fuzzy bag-of-words model for document representation**. *IEEE Transactions on Fuzzy Systems*, pages 1–1.
- Vitalii Zhelezniak, Dan Busbridge, April Shen, Samuel L. Smith, and Nils Y. Hammerla. 2018. **Decoding Decoders: Finding Optimal Representation Spaces for Unsupervised Similarity Tasks**. *CoRR*, abs/1805.03435.
- Vitalii Zhelezniak, Aleksandar Savkov, April Shen, and Nils Y. Hammerla. 2019a. **Correlation coefficients and semantic textual similarity**. In *NAACL-HLT*.
- Vitalii Zhelezniak, Aleksandar Savkov, April Shen, Francesco Moramarco, Jack Flann, and Nils Y. Hammerla. 2019b. **Don’t settle for average, go for the max: Fuzzy sets and max-pooled word vectors**. In *International Conference on Learning Representations*.