



HAL
open science

Unsupervised sentence representations as word information series: Revisiting TF–IDF

Juan-Manuel Torres-Moreno, Ignacio Arroyo-Fernandez, Gerardo Sierra, Grigori Sidorov, Carlos Francisco Mendez Cruz

► **To cite this version:**

Juan-Manuel Torres-Moreno, Ignacio Arroyo-Fernandez, Gerardo Sierra, Grigori Sidorov, Carlos Francisco Mendez Cruz. Unsupervised sentence representations as word information series: Revisiting TF–IDF. *Computer Speech and Language*, 2019, 56, pp.107-129. 10.1016/j.csl.2019.01.005 . hal-02550939

HAL Id: hal-02550939

<https://hal.science/hal-02550939>

Submitted on 20 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Unsupervised Sentence Representations as Word Information Series: Revisiting TF-IDF

Ignacio Arroyo-Fernández^{2,a,*}, Carlos-Francisco Méndez-Cruz^{1,b}, Gerardo Sierra^{2,c}, Juan-Manuel Torres-Moreno^{3,d}, Grigori Sidorov^{4,e}

^a*Universidad Nacional Autónoma de México (UNAM)*

^b*Centro de ciencias genómicas (CCG-UNAM)*

^c*Instituto de ingeniería (IIngen-UNAM)*

^d*Laboratoire Informatique d'Avignon (LIA-UAPV)*

^e*Centro de investigación en computación (CIC-IPN)*

Abstract

Sentence representation at the semantic level is a challenging task for Natural Language Processing and Artificial Intelligence. Despite the advances in word embeddings (i.e. word vector representations), capturing sentence meaning is an open question due to complexities of semantic interactions among words. In this paper, we present an embedding method, which is aimed at learning unsupervised sentence representations from unlabeled text. We propose an unsupervised method that models a sentence as a weighted series of word embeddings. The weights of the word embeddings are fitted by using Shannon's word entropies provided by the Term Frequency-Inverse Document Frequency (TF-IDF) transform. The hyperparameters of the model can be selected according to the properties of data (e.g. sentence length and textual gender). Hyperparameter selection involves word embedding methods and dimensionalities, as well as weighting schemata. Our method offers advantages over existing methods: identifiable modules, short-term training, online inference of (unseen) sentence representations, as well as independence from domain, external knowledge and language resources. Results showed that our model outperformed the state of the art in well-known Semantic Textual Similarity (STS) benchmarks. Moreover, our model reached state-of-the-art performance when compared to supervised and knowledge-based STS systems.

*Corresponding author

Email addresses: iaf@ciencias.unam.mx (Ignacio Arroyo-Fernández), cmendezc@ccg.unam.mx (Carlos-Francisco Méndez-Cruz), gsierram@ii.unam.mx (Gerardo Sierra), juan-manuel.torres@univ-avignon.fr (Juan-Manuel Torres-Moreno), sidorov@cic.ipn.mx (Grigori Sidorov)

¹Av. Universidad s/n Col. Chamilpa 62210, Cuernavaca, Morelos

²AV. Universidad No. 3000, Ciudad universitaria, Coyoacán 04510, Ciudad de México

³Université d'Avignon et des Pays de Vaucluse. 339 chemin des Meinajaries 84911, Avignon cedex 9, France

⁴Instituto Politécnico Nacional. Av. Juan de Dios Bátiz, Esq. Miguel Othón de Mendizábal, Col. Nueva Industrial Vallejo, Gustavo A. Madero 07738, Ciudad de México

Keywords: Sentence Representation, Sentence Embedding, Word Embedding, Information Entropy, TF-IDF, Natural Language Processing
2010 MSC: 00-01, 99-00

1. Introduction

Nowadays, the growth of information in digital media encourages the analysis of large amounts of text data. This is attracting attention from Data Science and Artificial Intelligence researchers, as well as from the Internet industry. Internet users are responsible for a meaningful part of this growth. They enter information into the network which is also leveraged for sharing knowledge. An important part of this knowledge is found at repositories such as question & answer forums, digital newspapers and digital encyclopedias.

Due to the innumerable duplication of the information at these repositories, several concerns arise as to the way users feed and consume knowledge. Some of these concerns include removing redundancies in question-answering forums or exploiting redundancies to assess the confidence of news in media or simply to compress text size. The accomplishment of this massive information processing is clearly infeasible for human reviewers. In this scenario, Statistical Natural Language Processing (NLP) methods are a substantial aid.

An approach to address these issues is to perform massive comparisons by considering the content of sentences or short snippets of text. These comparisons can be done by means of Semantic Textual Similarity (STS) systems (Hatzivassiloglou et al., 1999; Agirre et al., 2012). An STS system computes a similarity score (a real value) between a pair of sentences. This score indicates how similar the sentences of the pair are. Most STS systems incorporate a number of supervisory signals such as Knowledge Bases, encyclopedias, language resources (e.g. thesaurus and linguistic taggers built on the basis of a Part-of-Speech [PoS] tagger) and even similarity labels (Mihalcea et al., 2006). Nevertheless, for specialized texts (or for low-resourced languages) those resources are not available or are scarce. Furthermore, the scope of such systems is limited exclusively to the task of measuring textual similarity. Such a limitation generally obviates the step of vector representation of sentences. This is not desirable when we want to study statistical behavior of sentence meaning.

Another approach is the one followed in this work. It consists of embedding sentences (or short text snippets) onto vector spaces such that approximations to their meanings can be represented geometrically, i.e. sentence embeddings or sentence representations (Hinton et al., 1986; Elman, 1991). The main advantage of such an approach is that there exists the possibility of studying the statistical behavior of sentence meaning. As an additional and important benefit, sentence embeddings make it possible to leverage a number of NLP tasks, such as sentence clustering, text summarization (Zhang et al., 2012; Arroyo-Fernández, 2015; Arroyo-Fernández et al., 2016; Yu et al., 2017), sentence classification (Kalchbrenner et al., 2014; Chen et al., 2017; Er et al., 2016), paraphrase identification (Yin and Schütze, 2015), semantic similarity/

relatedness and sentiment classification (Arroyo-Fernández and Meza Ruiz, 2017; Chen et al., 2017; De Boom et al., 2016; Kalchbrenner et al., 2014; Onan et al., 2017; Yazdani and Popescu-Belis, 2013).

The usefulness of vector representation methods mainly depends on the characteristics of the text to be embedded into vector spaces (Salton and Buckley, 1988). On the one hand, most embedding methods provide well-suited representations of the content of texts for which individual size is on a relatively large scale (Salton et al., 1983; Martin and Berry, 2007; Le and Mikolov, 2014). For instance, book-sized texts (e.g. documents with hundreds of thousands of words or larger) are well represented by the importance of the words they contain (Spärk Jones, 1972). Accordingly, the size of these text objects (as a Bag-of-Words, BoW) suggests that by representing them we can satisfy shallow information necessities limited to the gist (topics) of the documents (Manning et al., 2009; Kintsch and Mangalath, 2011), e.g. Information Retrieval and document classification.

On the other hand, words are at the bottom end of the text size scale. At the word level, information necessities can be very general. That is, word representation methods could be components of practically any NLP system. These methods are mainly based on a general principle called *distributional hypothesis*, which states that similar words are used in similar contexts (Firth, 1957; Harris, 1968). In the NLP area, this linguistic principle is usually implemented as statistical estimates of word co-occurrence, i.e. word embedding methods. These statistical estimates provide word embeddings performing well enough in general purpose NLP applications (Baroni and Lenci, 2010; Mikolov et al., 2013a; Pennington et al., 2014; Baroni et al., 2014; Bojanowski et al., 2016).

The problem of modeling sentences is still open. For the cases of documents or words, most applications expect representations encoding text content or word use. Nonetheless, in the case of sentences, application users can expect composite representations providing much more specific information, e.g. what is declared or denied about something (Pereira, 2000; Meza-Ruiz and Riedel, 2009; Collobert et al., 2011; Kintsch and Mangalath, 2011). Thus, state-of-the-art sentence representation methods can be highly dependent on the application and on its specificity. So it is difficult to keep their performance and behavior uniform/stable in several scenarios (Pham et al., 2015; Pagliardini et al., 2017). Advancing the state of the art on sentence representation can be specially useful when only unlabelled text is available for learning sentence representations or even for applications to low-resourced languages.

In this work we address the problem of sentence representation by means of the following hypothesis. It is possible to represent sentences well by exploiting the link between the contexts learned by word embeddings and the entropy of such embedded words. In order to confirm our hypothesis, we present a modular model that consists of a weighted series of pretrained word embeddings. In this framework, the weights of the word embeddings are fitted by using unsupervised learning based on Shannon’s entropy (Shannon, 1949), i.e. TF-IDF (Term Frequency–Inverse Document Frequency). Thus, we take into account word importance, both at corpus level and at sentence level.

We evaluated our model, called *Word Information Series for Sentence Embedding (WISSE)*, by using well-known STS benchmarks provided by the SemEval competition (i.e. SICK [Sentences Involving Compositional Knowledge, 2014] and SemEval [2016]). Our results showed that WISSE outperformed (or was comparable to) strong state-of-the-art methods in such benchmarks. Additional advantages were observed, which were mainly due to the modularity and low computational cost of our model: short-term training, independence from domain, external knowledge and language resources, as well as online inference of (unseen) sentence representations.

The rest of this paper is organized as follows: Section 2 presents the related work. Section 3 presents the main differences between STS systems and sentence representation methods. Section 4 exposes the motivations for our method. Section 5 presents the modules composing our model. In Section 6 we explain the constitution of our model and how the modules composing it interact. In Section 7 we explain the design of our experiments and their objectives. Section 8 presents the obtained results and Section 9 addresses the discussion about such results. Section 10 provides insights on possible improvements, including advantages and disadvantages. Finally, in Section 11 the conclusions derived from this work are presented.

2. Related work

The method proposed in this paper simultaneously falls into two categories of sentence representation methods, whose related work is reported in this section. The first category deals with models and methods that unsupervisedly embed short text snippets or sentences. Such an embedding is performed directly from unlabeled (plain) text data into vector spaces. The main methods are based on statistics and neural networks. The second category deals with models that use any form of weighted sum of word embeddings. This mechanism embeds text snippets or sentences into vector spaces by using pretrained word embeddings. In this framework the weights of the word embeddings can be fitted with both supervised and unsupervised learning.

2.1. Unsupervised methods for sentence representation

In this subsection we briefly describe a number of unsupervised methods whose aim is to build sentence representations for general purposes. These methods do not use external resources or supervisory signals, some of the main traits of the model presented in this paper.

2.1.1. Statistical methods

A popular statistical representation method was originally used in Information Retrieval applications, by which documents were represented. The representation consists of the TF-IDF transform of document vectors such that their components are, mainly, word frequencies. Currently there are multiple heuristics for computing the components, e.g. word presence/absence (binary),

smoothed logarithm, etc. (Salton et al., 1983; Salton and Buckley, 1988). The transform gives a *term-document* matrix as a result. Due to its monotonic nature, this model assigns relatively high weights to rare (infrequent) words and relatively low weights to very frequent words (Spärk Jones, 1972). According to information theory, this schema weighs the information conveyed by each word of the vocabulary (Aizawa, 2003; Robertson, 2004). The weighting method consists of a logarithmic rescaling of the frequency of each word within a document. At the end, this logarithm linearizes the exponential distribution of word types through the corpus. In our experiments, we included BoW for sentence representation as a baseline method.

Latent Semantic Analysis (LSA) takes as its input a term-document matrix created by means of the BoW method (Landauer et al., 1998). The sparse vectors of this matrix are transformed into document vectors, the of which are the projection weights to a user-defined number of eigenvectors of the term-document matrix. The transformation is computed by means of the Singular Value Decomposition (SVD) method. The number of eigenvectors is associated with the number of topics supposed to be present in the collection of documents (Martin and Berry, 2007).

2.1.2. Neural sentence representations

The method called *Doc2Vec* uses a neural network to build sentence/paragraph vectors that can be used for general purposes (Le and Mikolov, 2014). This method uses word embeddings previously learned from fixed-length segments of text (sliding word windows). The words of a sentence are associated with the corresponding word embeddings (Mikolov et al., 2013b). These embeddings then are used as evidence to predict a virtual word embedding which does not represent a word, but a sentence instead.

As an extension of *Doc2Vec*, the neural model proposed by Kiros et al. (2015) produces sentence embeddings from the hidden states of a Recurrent Neural Network (RNN). In this framework (the Skip-thought vectors), the two sentences surrounding a center sentence are a context window. The RNN maps these sentence contexts to its last hidden state, which is taken as a sentence embedding.

The architecture called *FastSent*, proposed by Hill et al. (2016), is based on the Glove model (Pennington et al., 2014) (see Section 5.1). Furthermore, the authors instantiate the architecture of Skip-Thought vectors (Kiros et al., 2015). This combined network uses a precomputed matrix that merges co-occurrence information from both words and sentences. Additionally, Hill et al. (2016) addresses the STS problem as one of machine translation. This helps to learn sentence representations by simulating negative examples to a Sequential (Denoising) Auto Encoder, S(D)AE, which uses such examples to learn a negative model. Thus, unseen sentences can be built in opposition to the jointly learned negative (adversarial) model.

The work of Wieting et al. (2016) proposes a meaningful difference with respect to most word embedding methods. This model, called *CHARAGRAM*

learns embeddings of character n -grams (Bojanowski et al., 2016). Character embeddings are simply averaged in order to compose words. Actually the same operation is performed when sentence embeddings are needed, i.e. the obtained word embeddings are averaged to obtain sentence representations in the *CHARAGRAM-PHRASE* model.

The neural sentence representation model called C-PHRASE relies on dependency/constituency parsing (Bentivogli et al., 2016). The idea is very similar to the one proposed by Levy and Goldberg (2014) for word embeddings (Section 5.1.2). That is, the word co-occurrence is constrained by structure dependencies rather than by word-context windows scanning the input corpus.

The model called *Sent2Vec* is similar to that proposed by Le and Mikolov (2014). The authors extended Doc2Vec for considering sentences, instead of fixed-length context windows. Additionally, this model can consider word n -grams or even the dynamic length of the context window as a modification of the subsampling approach proposed by Mikolov et al. (2013a). This model is also very similar to that proposed by Bojanowski et al. (2016) for word embeddings (Section 5.1.3). That is, the architecture can learn a distribution of labels for a given training example.

Although sentence representation methods based purely on deep learning have shown competitive performance based on a number of benchmarks, their computational cost can be a significant bottleneck. Some of these methods need large amounts of data and even weeks of training on GPUs to perform reasonably well (Kiros et al., 2015).

Most state-of-the-art methods are purely neural network based. Unlike such an approach, our method uses neural models only for training word embeddings (our word embedding module).

2.2. Weighted sum of word embedding methods

The notion of weighting word embeddings composing sentences was introduced by Ji and Eisenstein (2013). A recent extension of such work is presented in (Yin and Schütze, 2015). The authors appeal to supervised learning of weights for each word within a sentence. Their approach focuses on improving precision in paraphrase identification. This method rescales (enforced/penalized) the IDF of shared words of a pair of sentences. Herein, two complementary distributions of events are considered. First, the probability distribution that a given word taking place in sentences that are paraphrases. Second, the probability distribution that a given word taking place in sentences that are not paraphrases. The *Kullback-Leibler* divergence of both distributions is computed to enforce IDF weights of shared word embeddings between paraphrases. The authors use negative factorization of the rescaled sentence matrices in order to build vector features for paraphrase classification.

In (Zheng and Callan, 2015; Brokos et al., 2016), the authors improve Information Retrieval applications by using word embedding weighting. The embeddings of documents and queries are weighted via IDF. The goal is to compute the average similarity between query term embeddings and word embeddings of

documents for retrieval. In (Kenter and de Rijke, 2015) a similar approach is proposed. From partial measurements, a regression function is estimated from a training dataset in order to predict the rescaling of the word embeddings of a test set.

In (De Boom et al., 2016), the authors propose a supervised approach for learning the word embedding weights for text snippets (short texts). Their method represents snippets of paragraph-like lengths. A binary classification problem (relatedness/unrelatedness) is posed to learn the required weights. In this method, the IDF weights are used to rank the order of the weighted embedding summation. The authors show that unimportant words induce too much bias to semantic similarity measurement⁵. Therefore their model drops unimportant words from the sentence representation. Thus their model learns uniquely the weights of surviving word embeddings. Next, the weighted word embeddings are averaged in order to obtain the final snippet representation. The authors pointed out that such an average works well for texts of about 30 words in length. Nonetheless, they also propose additional modifications to their model for variable text length. Unlike this approach, our study is for sentence representation.

Another related method is the one recently proposed by Ferrero et al. (2017). Their model uses PoS tag weights. These weights are merged to the associated IDF weights from the analyzed dataset of the STS sentence pairs. Both the PoS weights β_{POS} and the IDF weights β_{idf} of each word w are combined using a product of powers. By means of such a combination, a unique word weight ϕ_w is obtained, i.e. $\phi_w = \beta_{POS}^\alpha \beta_{idf}^{1-\alpha}$. As both β_{POS} and α are free parameters, the model learns them in a supervised fashion by using manually annotated semantic similarity scores. Thus, both word importance and word PoS tags contribute to weight word embeddings in order to build a sentence representation.

In (Arora et al., 2017) the authors use word embedding weighting in both supervised and unsupervised approaches. The weights were learned as part of a multinomial distribution. This multinomial is parameterized according to the probability of a word appearing along with other words (i.e. the words of a sentence). This idea is actually an extrapolation of the principles behind distributed representations for word embeddings (Mikolov et al., 2013b). However, this model additionally considers a balance (a linear convex compensation) between the probability of a word to occur within a *discourse* and its probability to occur within a sentence.

Most methods we presented in this subsection are supervised. An exception is the unsupervised method proposed by Arora et al. (2017), which is the most similar to ours. This method performs unsupervised learning of co-occurrence weights. Unlike to such an approach, ours performs unsupervised learning of weights based on information amount (which constitutes our information-

⁵We think that this noisy behavior is not necessarily due to the property of a word of being noisy, but more probably due to the over/underrepresentation issues in word embedding methods.

theoretic module).

3. STS and the distinction between STS systems and sentence representation methods

Since we evaluated our sentence representation method by means of STS tasks, we briefly explore such tasks' context in this section. Furthermore, we show a comparison between sentence representation methods (like the proposed one in this paper) and STS systems.

There are a number of proposed methods to assess semantic similarity between pairs of sentences. Most of these proposals emerge from benchmarks like the one used in the SemEval STS competition (Agirre et al., 2012). Given a pair of sentences, the aim of the STS task is to determine a similarity score (a real value). This score indicates just how similar the sentences of the pair are. The higher the score is, the higher the measured semantic similarity.

STS benchmarks evaluate the correlation coefficient between the similarity assessed by some algorithm and the similarity assessed manually by humans. This coefficient (the Pearson's coefficient) is a real number $\rho(\cdot, \cdot) \in [-1, 1]$. For instance, let $d = \{d(S_a^{(1)}, S_b^{(1)}), \dots, d(S_a^{(\ell)}, S_b^{(\ell)})\}$ be the similarities of a set of ℓ pairs of sentences. Also let $y = \{y_1, \dots, y_\ell\}$ be the gold standard of similarities manually annotated. In this example as $\rho(d, y)$ approaches 1.0, it means that the STS algorithm producing d performs well on the benchmark.

STS systems attaining relatively high performances have recently been developed. These systems correlate about 80% with similarity scores annotated by humans (Cer et al., 2017). Most systems integrate combinations of multiple algorithms providing partial scores from a number of aspects of the sentences (Sultan et al., 2014). For instance, a typical similarity score between sentences of a pair can be obtained as follows (Pilehvar and Navigli, 2015; Brychcm and Svoboda, 2016):

$$d(S_a, S_b) = \alpha_1 d_1 + \dots + \alpha_n d_n.$$

In the $d(\cdot, \cdot)$ score, the α_i s are weights controlling the influence of each aspect or attribute shared between sentences S_a, S_b . Such weights may be either user-defined or learned in a supervised fashion. Each d_i is the overlapping score of the i th attribute (or category of attributes) aligned between sentences (e.g. words aligned by syntactic category, words aligned by their vicinity within a semantic graph, etc.). Notice that each $\alpha_i d_i$ is actually a partial similarity score contributing independently to $d(\cdot, \cdot)$.

There are also semisupervised systems whose main feature is the use of external resources such as knowledge bases, thesauruses and dictionaries (Pilehvar and Navigli, 2015; Kenter and de Rijke, 2015; Brychcm and Svoboda, 2016). Other methods also incorporate supervisory signals provided by PoS taggers, dependency/semantic parsers and Neural Networks (Rychalska et al., 2016).

In our case, we evaluate the performance of our sentence representation method by using popular STS benchmarks, i.e. SemEval (Agirre et al., 2016)

and SICK (Bentivogli et al., 2016). In this context, the evaluation of any sentence representation method like the ones shown in Section 2 is very similar with respect to the evaluation of STS systems. Nonetheless, the interpretation of the results is quite different.

When evaluating sentence representations, we only define a similarity function as an algorithm for comparing vector representations. For instance, the cosine similarity: $\cos(\theta) = \hat{y}(s_a, s_b) = (s_a \cdot s_b) / (\|s_a\| \cdot \|s_b\|)$. In this scenario, we should be focused on details involved in the construction of the representations. Thus, a set of similarities $\hat{y} = \{\hat{y}(s_a^{(1)}, s_b^{(1)}), \dots, \hat{y}(s_a^{(\ell)}, s_b^{(\ell)})\}$ is obtained from comparing a set of ℓ pairs of sentence representations.

Let the gold standard of similarities manually annotated be $y = \{y_1, \dots, y_\ell\}$. As a good result on the STS benchmark we expect that the correlation coefficient to approach one, i.e. $\rho(\hat{y}, y) \rightarrow 1.0$. It means that our sentence representation method generating the embeddings $s_{(\cdot)} \in \mathbb{R}^d$ performs well on the benchmark and with respect to the cosine similarity algorithm. Therefore, such representations encode a good approximation of the human sentence-meaning criterion given by y .

In general, STS systems do not need to represent sentences at all. In fact, there is a remarkable distinction between the aim of an STS system and the aim of a sentence representation method evaluated in STS tasks. That is, most STS systems are designed to beat the STS ranking, while sentence embedding methods are designed to provide distributional semantic representations of sentences. See Table 1.

Table 1: Comparison of characteristics between STS systems (**sts**) and sentence representation methods (**repr**). Column names: Access to sentence representations (**Access**={Yes/No}), Partial similarity scores (**Partial**={Yes/No}), use of external (knowledge) resources (**Resources**).

System/method	sts/ repr.	Parsing	Resources	Partial	Access
Rychalska et al. (2016)	sts	Dependency	WordNet	Supervised	No
Han et al. (2013)	sts	Chunking	WordNet	Yes	No
Sultan et al. (2014)	sts	Chunking, named entities, dependency	WordNet	Yes	No
UWBunsup (Brychcin and Svoboda, 2016)	sts	Chunking	No	No	No
LSA	repr.	No	No	No	Yes
BoW	repr.	No	No	No	Yes
Sent2Vec	repr.	No	no	No	Yes
Doc2Vec	repr.	No	No	No	Yes
FastSent	repr.	No	No	No	Yes
Skip-Thoughts	repr.	No	No	No	Yes
Glove-WR	repr.	No	No	No	Yes
C-PHRASE	repr.	Constituent structure	No	No	Yes
WISSE	repr.	No	No	No	Yes
WISSE + dependency based word embeddings	repr.	Dependency	No	No	Yes

In this paper, we focused our comparisons on unsupervised learning methods

which require neither labeled data nor external resources. Moreover, notice that we focused on sentence representation methods rather than on STS systems.

4. Motivation for weighted series of word embeddings

Currently there are multiple approaches for building sentence representations in vector spaces. Nonetheless, the complexity of the problem of keeping stable sentence representations through multiple scenarios remains a bottleneck in the NLP area. This is mainly because there exist linguistic and knowledge resources unique for a limited subset of information necessities and languages. So research on efficient unsupervised sentence representation methods can offer a promising approach to such a limitation. In this section we present the intuitions that motivated our contribution. These intuitions are given mainly in the sense of orthogonality of word embeddings as a desired trait of weighted series.

4.1. Composition in distributional semantics

Mitchell and Lapata (2010) propose a number of candidate models for semantic composition that are empirically tested as heuristics yielding promising results. Among the candidate models, the asymmetric composition is particularly interesting for us. This model is a weighted sum of word embeddings. Its purpose is to approximate the meaning compositionality of short phrases (e.g. “random variable”, “meaning composition”). In this framework, the asymmetry is posed as the linguistic feature such that the *head*[H] of a phrase must be more important than the *dependent modifier*[MD]. See examples (1) and (2):

- (1) *random*[DM] *variable*[H]
- (2) *meaning*[DM] *composition*[H].

The asymmetric composition model is given by:

$$p = \alpha x_{[\text{DM}]} + \beta x_{[\text{H}]}, \tag{1}$$

where it must be verified that $\alpha < \beta$. This inequality reflects the difference between the importance of the constituents (i.e. the word embeddings $x_{[\text{DM}]}, x_{[\text{H}]} \in \mathbb{R}^n$) of the resulting phrase $p \in \mathbb{R}^n$. According to Tian et al. (2017), coefficients α, β are scalars drawn from a monotonic function. In this work, we consider that a reasonable choice for such a monotonic function is the Shannon’s entropy (Shannon, 1949; Charniak, 1996; Aizawa, 2003). The asymmetric model takes into account both word order and linguistic features determining the syntactic category of the resulting phrase embedding p .

Since the composition approach (1) is plain and natural, it has encouraged recent work. For instance, Tian et al. (2017) described theoretical conditions for the vector averaging operation as a model of composition in distributional semantics. Given word embeddings $x_{[\text{DM}]}$ and $x_{[\text{H}]}$ that are geometrically uncorrelated, then making $\alpha = \beta = \frac{1}{2}$ actually causes p to approach zero. In other words, the average operation causes embeddings of words co-occurring with low

or moderated frequencies to cancel each other, so $p \rightarrow 0$. This effect suggests a linguistic intuition: uncorrelatedness (including orthogonality) between the meanings of words of a phrase occurs when the speaker constructs composite meanings. In contrast, as word meanings are correlated in context (e.g. words co-occurring frequently, like "cell phone"). Frequent co-occurrences lead to simple meanings that can easily be represented by means of the average between word embeddings. In this case, p is actually encoded as an implicit embedding shared by $x_{[\text{DM}]}$ and $x_{[\text{H}]}$. From our experiments, we interpret that the aforementioned observations explain the low performance of the simple average of word embeddings for representing whole sentences (Section 8).

4.2. The sparseness in neural language models

Word embedding models have one characteristic in common: the sparseness of word co-occurrence statistics induces orthogonality (Elman, 1991). Now, we present the main cases during the training of word embeddings where orthogonality is relaxed and where it approximately holds. For instance in recent work on neural models (Bengio et al., 2003; Mikolov et al., 2013a), binary sparse vectors are used for representing input words as categorical variables (one-hot encoded vectors). These vectors build a canonical basis for $\mathbb{R}^{|V|}$. This basis encodes the vocabulary as an orthonormal set

$$e = \{e_1, \dots, e_{|V|}\} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ \vdots & \ddots & \dots & \vdots \\ 0 & \dots & 0 & 1 \end{pmatrix}$$

before training (Elman, 1991). After training, the projection layer of the neural model has learned a transformation relaxing the initial orthogonality of words co-occurring frequently. Conversely, for words that do not co-occur frequently, orthogonality is impregnated or imposed onto their associated embeddings.

4.3. The extreme values of training on co-occurrences

Let us consider the case of Glove (Pennington et al., 2014). Suppose that the word w_i does not (or almost does not) co-occur along with the context $c_j = \{w_1, \dots, w_{i+r}\}$. Also, let x_i and φ_j be the corresponding embeddings. Such embeddings are learned by the Glove's loss:

$$\mathcal{J}(x_i, \varphi_j) = \sum_{i,j}^V f(c_{ij}) (\langle x_i, \varphi_j \rangle - \log c_{ij})^2. \quad (2)$$

If the word w_i does not co-occur within the context c_j , then $c_{ij} \rightarrow 1$ and therefore $\log c_{ij} \approx 0$. As the objective (2) requires that $(\langle x_i, \varphi_j \rangle - 0)^2 \rightarrow 0$, then the dot product $\langle x_i, \varphi_j \rangle \approx 0$. This implies that the word embeddings tend to be orthogonal.

A different tendency can be observed when w_i co-occurs frequently along with the context c_j . The training leads to $\log c_{ij} > b$, so that $\langle x_i, \varphi_j \rangle \approx b$ for a

sufficiently large $b > 0$. Therefore the word embeddings are linearly dependent proportionally to b .

Notice that this analysis considers the extreme values of co-occurrence. Depending on the stability of the embedding algorithms, orthogonality is uniquely a tendency that approximates an ideal distribution.

4.4. Merging

As an extension of the geometrical observations on compositionality and orthogonality, we propose weighted series of word embeddings for sentence representation. Weighted series can be seen as the decomposition of a sentence into basis vectors. Mathematically, a basis is an orthogonal set spanning a semantic space of sentences (Fourier, 1822). As being convenient for our proposed model, we decided to relax this concept. In our weighted series we include word embeddings just as produced by their algorithms (e.g. Word2Vec and Glove).

Pennington et al. (2014) offers a proof that the distribution of co-occurrences is bounded by a *generalized harmonic number* (i.e. a set of logarithmic distributions). Such a proof and the analysis presented in sections 4.1, 4.2 and 4.3 suggest that the angles between word embeddings are also logarithmically distributed as they are combined to transmit information. Thus, such a distribution relates Shannon’s entropy, co-occurrence statistics and orthogonality. This reasoning in turn suggests that a weighted series model admits, on the one hand, the weighted contribution of word embeddings encoding composite meanings. On the other hand, this model also admits the weighted contribution of lexicalized phrases whose constituent words share implicit embeddings. Thus, theoretically a weighted series of word embeddings is built by two main subsets of word embeddings. The former subset contains a few uncorrelated (probably orthogonal) word embeddings carrying specific information. Also, the second (disjoint) subset contains correlated word embeddings carrying almost no information. Notice how this idea resembles the idea put forward in (1).

5. The modules of our sentence representation method

5.1. The word embeddings module

In this section we will describe the word embedding methods that WISSE employs as one of its modules. One of the most popular methods is termed Word2Vec (W2V) (Mikolov et al., 2013a), which is a neural language model inspired by ideas proposed previously (Hinton et al., 1986; Bengio et al., 2003). Recently other methods have been proposed and they are still growing in popularity, e.g. Glove (Pennington et al., 2014) and FastText (Bojanowski et al., 2016). These methods share similar elements in general. Their aim is to train a learning model whose hidden parameters are random vectors. The model learns from multinomial distributions modeling the co-occurrence of each word type along with other words. These other words are encapsulated within sliding windows (i.e. target contexts for each word type (Rong, 2014)). Once the model is trained, its hidden parameters are used as word embeddings.

5.1.1. Word2Vec

There are two possible neural architectures producing word embeddings with W2V: the Skip-gram and the CBoW. In this work, we used the Skip-gram neural architecture for training word embeddings.

Let $V = \{w_1, \dots, w_t, \dots, w_{|V|}\}$ be the vocabulary of a context corpus D . The main idea behind Skip-gram is to (supervisedly) classify all word types w_t as they are instantiated ($w_i = w_t$) though such a corpus $D = \{w_1, \dots, w_i, \dots, w_{|D|}\}$. Multiple target labels are simultaneously associated with each w_t . Each target set is the vocabulary of a sliding window having r labels being the words around w_i . Thus, the target associated with w_i takes the form $c_i = \{w_{i-r}, \dots, w_{i+r}\} \setminus w_i$. The parameters of the distribution modeling the set $\{c_1, \dots, c_i, \dots, c_{|D|}\}$ must be learned by the Skip-gram classifier.

Before training of the classifier, a preprocessing step transforms the input unlabeled text into a labeled one. This amounts to obtain a training set of the form $(w_1, c_1), \dots, (w_i, c_i), \dots, (w_{|D|}, c_{|D|})$. As Skip-gram is actually a Neural Network, its output layer can be written in terms of its hidden layers as follows:

$$P(c_i | w_i = w_t) = P(w_{i-r}, \dots, w_{i+r} | w_t) = \frac{\exp(x_t, \varphi_{c_i})}{\sum_{w \in V} \exp(x_t, \varphi_w)}. \quad (3)$$

The vectors $x_t \in \mathbb{R}^d$ are word embeddings of the word types w_t and they are actually the columns of the first hidden layer of the Skip-gram. The dimension of the word embeddings is d , which equals the number of hyperplanes described by the first hidden layer. Entries of matrix φ_{c_i} are the weights of the output layer, corresponding to the word types $w_{i-r}, \dots, w_{i+r} \in c_i$. Entries of matrix $\varphi_w |_{w \in V}$ are all weights of the output layer.

5.1.2. Dependency-based word embeddings (Dep2Vec)

One of the embedding methods we used as part of our experiments is an extension of W2V. The extension consists of a substantial modification of the notion of context window that uses dependency parsing as a surrogate for the standard co-occurrence window c_i in (3). Given a word w_i , dependency parsing is used for establishing context restrictions. The pair $(w_{i \pm r}, w_i)$ is considered to co-occur within the same context whenever its words are related by a grammatical dependency. For example, according to the mentioned restrictions, the co-occurrence of the words *John* and *sneezes* has much higher probability than the co-occurrence of the words *house* and *sneezes*. *House* and *sneezes* can co-occur within a vicinity of words, but they do not hold a grammatical dependency. Therefore this pair is not considered as a co-occurring one. This is because of the object-subject distributional dependency restricting the actions that people and houses can do, i.e. a people sneeze, but a houses do not (Baroni and Lenci, 2010; Levy and Goldberg, 2014; Pagliardini et al., 2017).

5.1.3. FastText

The FastText model is another extension of W2V (Bojanowski et al., 2016). This model is a combination of the CBoW and Skip-gram architectures. Fast-

Text is designed to segment the input text into contiguous n -grams of characters g_i . These contiguous n -grams are then grouped into context windows c_i . As in the case of W2V, c_i is associated with w_i and $w_i \notin c_i$. Thus, given a context window $c_i = \{g_1, \dots, g_{|c_i|}\}$ the model should predict a target window $c'_i = \{g_1, \dots, g_{|c_i|}, w_i\}$, i.e. $P(c'_i|c_i) = P(g_1, \dots, g_{|c_i|}, w_i|g_1, \dots, g_{|c_i|})$. Notice that both c_i and c'_i are a bit different than c_i built in (3). Furthermore, the target window contains the word w_i which is composed by the n -grams g_i . This allows the model for different word embeddings for the word *as* and for the bigram **as** composing the word *whereas*. The same mechanism allows FastText to infer out-of-vocabulary word embeddings.

5.1.4. Global vectors for word representation (Glove)

Global vectors for word representation (Glove) is a bit different from most of its counterparts. By bringing back pioneer methods like LSA, this model transforms a sparse matrix of word co-occurrences into co-occurrence probabilities (Pennington et al., 2014). The matrix of co-occurrences C has entries c_{ij} and each of them is the probability that the word w_i co-occurs along with other words w_j within a context window. The model estimates a regression function on the non-zero log probabilities. This considers as independent variables the set of word embeddings $x_i \in \mathbb{R}^d$.

$$\mathcal{J}(x_i, \varphi_j) = \sum_{i,j}^V f(c_{ij})(\langle x_i, \varphi_j \rangle - \log c_{ij})^2, \quad (4)$$

where x_i corresponds to the word w_i and $\varphi_j \in \mathbb{R}^d$ is the embedding of the context window spanning over other words w_j . The weighting function $f(c_{ij}) = \left(\frac{c_{ij}}{\max\{c_{ij}\}}\right)^\alpha$ allows \mathcal{J} to be adapted to the importance of co-occurrence probabilities c_{ij} .⁶ Overall, the convex objective (4) is conceived to optimize the correspondence between the inner product $\langle x_i, \varphi_j \rangle$ and the log difference $\log c_{ij} = \log \frac{P(w_j|w_i)}{P(w_i)}$.

5.2. The information-theoretic module

TF-IDF is a normalized log transformation providing three sources of information about the importance of words (Spärk Jones, 1972; Salton et al., 1983; Aizawa, 2003). In our proposed model, these sources of information are merged so as to build a contribution function based on Shannon’s entropy.

Given a vocabulary V , the first source of information is the entropy of a sentence $s_j \in S$ due to the set of sentences S (the corpus). The second source is the entropy of each word $w_i \in V$ due to the subset of sentences sharing it. The third source is the decrease of entropy of S as w_i is observed in a particular sentence $s_j \in S$.

⁶Similar considerations are also referred to as “sub-sampling” in Mikolov et al. (2013a) and Bojanowski et al. (2016).

Formally, according to Aizawa (2003), a sentence s_j has probability $P(s_j) = 1/N_S$ and all $s_j \in S$ are equally probable, so they are uniformly distributed. Now, it is defined the first information source as the entropy of each sentence s_j due to S :

$$H(s_j) = - \sum_{s_j \in S} P(s_j) \log P(s_j) \approx - \log \frac{1}{N_S}, \quad (5)$$

where N_S is the total number of sentences, i.e. $j = 1, 2, \dots, N_S$.

In a similar way, we can define our second source of information. It is the entropy of a word w_i due to the subset $\zeta_i \subset S$ of sentences sharing such a word:

$$H(S|w_i) = - \sum_{s_j \in S} P(s_j|w_i) \log P(s_j|w_i) \approx - \log \frac{1}{N_{w_i}}, \quad (6)$$

where N_{w_i} is the cardinality of ζ_i . One interpretation of the right side of (6) is that all elements of $\zeta = \{\zeta_i \subset S : i = 1, 2, \dots, |V|\}$ are equally likely and uniformly distributed.

Prior to computing our third information source we need to compute a conditional information source provided the first (5) and the second (6) information sources. It is the decrease in the entropy of S and ζ_i given that w_i is observed (Osteyee and Good, 1974). Such a decrease is also interpreted as the expected mutual information gained by ζ_i and S as each w_i is observed:

$$\begin{aligned} I(V, S) &= \sum_{w_i \in V} P(w_i) [H(S) - H(S|w_i)] \\ &\approx \sum_{w_i \in V} \frac{f_{w_i}}{F} \left(\log \frac{1}{N_{w_i}} - \log \frac{1}{N_S} \right). \end{aligned} \quad (7)$$

The probability of any w_i is computed as $P(w_i) = f_{w_i}/F$. f_{w_i} is the frequency of w_i in S and F is the sum of frequencies of all words in S . The conditional information described by (7) is a computation on the whole corpus. The marginal information amount associated with each word given a particular sentence can be obtained from each term of the summation (7). Such a marginal quantity is our third information source and it is defined explicitly in Section 6.

6. The proposed sentence representation model

In this paper, we propose to use the information provided by the entropies of words in a corpus (TF-IDF) as a weighting approach for series of word embeddings of the form

$$s(x, \varphi) = \sum_{w_i \in s} \varphi_i x_{w_i}. \quad (8)$$

In the general model (8), the unknown coefficients φ_i weigh the contribution from each word embedding x_{w_i} to the sentence representation $s(x, \varphi)$. A natural instinct is to optimize $\varphi_1, \dots, \varphi_{|s|}$ without (or with few) assumptions about

them. However, in this work we propose taking into account information features of each word obtained from a corpus as prior knowledge (Schölkopf et al., 1997). This prior knowledge relies on the fact that topical meaningful words (e.g., noun phrases and named entities) have relatively low probability of being used (Robertson, 2004), therefore they are informative. Conversely, non-topical words (e.g., prepositions and determinants) have a high probability of being used. Therefore, they are much less informative. Note that there is a correlation between syntactic features and how informative a word is (Mitra et al., 1997; Pereira, 2000; Mitchell and Lapata, 2010; Badarınza et al., 2017).

Let us instantiate our information-theoretic module (Section 5.2):

$$I(V, S) \approx \sum_{w_i \in V} \frac{f_{w_i}}{F} \left(\log \frac{1}{N_{w_i}} - \log \frac{1}{N_S} \right). \quad (9)$$

Now we use the prior knowledge it acquired once it has been trained on the corpus S . This means that, given the decomposition form of such a model, we can take the TF and the IDF vectors from it separately. Therefore the IDF vector $\varphi_S \in \mathbb{R}^{|V|}$ is given by:

$$\varphi_S = \left(\log \frac{1}{N_{w_1}}, \dots, \log \frac{1}{N_{w_{|V|}}} \right) - \log \frac{1}{N_S}, \quad (10)$$

where N_{w_i} is the number of sentences sharing each w_i in (9) and N_S is the total number of sentences in S . Note that each component of φ_S gives marginally the first and the second information sources given in Section 5.2.

Now, from each term $P(w_i) = \sum_{s_j \in S} P(w_i|s_j) \approx \sum_{s_j \in S} f_{ij}/F$ in (9), we define f_{ij} as the frequency of w_i occurring within some sentence s_j . Therefore the probability of w_i to occur in s_j is given by $P(w_i|s_j) \approx f_{ij}/F$. Thus, we define our third information source. As a particular w_i is observed in a particular s_j , it constitutes a conditional event decreasing the entropy of S (which also increases its informativeness). Each term of (9) therefore provides the informativeness of each particular event (w_i, s_j) , with which we define the TF vector:

$$\varphi_{w_i}^{(s_j)} = \left(0, \dots, \frac{f_{ij}}{F}, \dots, 0 \right) \quad (11)$$

For consistency with φ_S , the TF vector $\varphi_{w_i}^{(s_j)} \in \mathbb{R}^{|V|}$ has a unique nonzero component i in (11). Now we merge our three information sources into the scalar weights $\varphi_i = \langle \varphi_{w_i}^{(s_j)}, \varphi_S \rangle$ for each word embedding x_{w_i} of the series (8), i.e. the associated TF-IDF weights.

Suppose that any of the models described in Section 5.1 is already trained, e.g. via Glove (4), so we can now instantiate our word embedding module. Let us pick up any word embedding x_{w_i} from it.

Let a sentence representation $s(\cdot, \cdot)$ be a vector derived from a textual construction whose elements (word embeddings) “live” in a semantic space $X \subset \mathbb{R}^d$.

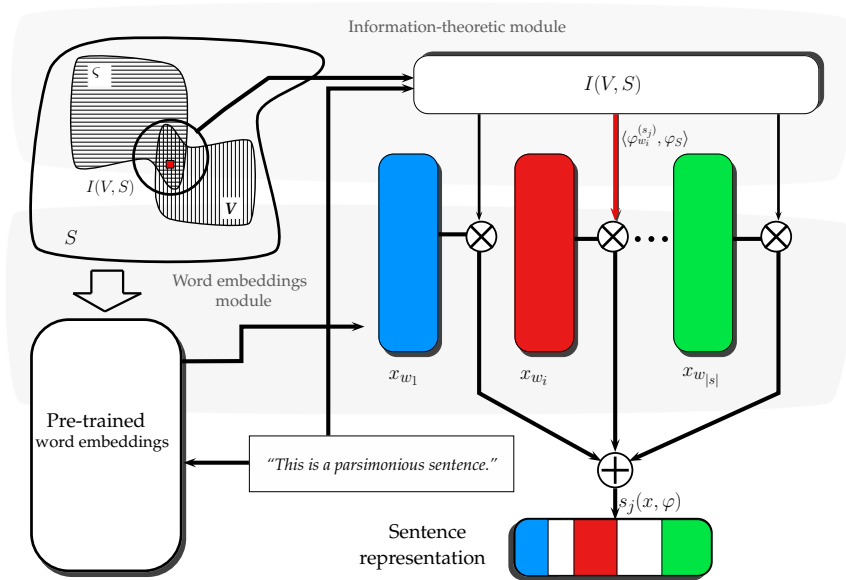


Figure 1: Given a sentence (e.g. “*This is a parsimonious sentence.*”), our model (12) asks to the word embedding module the word embeddings $\{x_{w_1}, \dots, x_{w_{|s|}}\}$ associated with the words constituting such a sentence. In turn, the model asks to the information-theoretic module for the IDF vectors φ_S and for the TF vectors $\varphi_{w_i}^{(s_j)}$ associated with each of these words. These vectors are multiplied, which gives us the TF-IDF weights $\langle \varphi_{w_i}^{(s_j)}, \varphi_S \rangle$ for each word in the sentence. Once the needed TF-IDF values are obtained, the model assigns them as weights to their associated word embeddings. According to our model, the sentence representations $s_j(\varphi, x)$ are obtained simply by summing the weighted word embeddings.

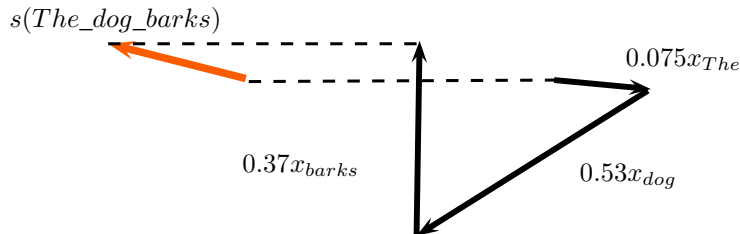


Figure 2: A **hypothetical** sketch of the sentence vectors for the sentence “*The dog barks*” and the actual weights computed from the corpus. The vector projected by the dotted lines (left side) is the hypothetical summation of constituents: the sentence representation.

Now let $\varphi = (\varphi_1, \dots, \varphi_{|s_j|}) = (\langle \varphi_{w_1}^{(s_j)}, \varphi_S \rangle, \dots, \langle \varphi_{w_{|s_j|}}^{(s_j)}, \varphi_S \rangle) \in \mathbb{R}^{|s_j|}$ be the coefficients of the weighted series (8), which have been learned by (9) in an unsupervised fashion. Hence, we rewrite (8) as follows:

$$s_j(x, \varphi) = \sum_{w_i \in s_j} \langle \varphi_{w_i}^{(s_j)}, \varphi_S \rangle x_{w_i}. \quad (12)$$

As a sentence representation $s_j(x, \varphi)$ is built, the learned vector φ regulates the amount of information provided by each constituent word embedded in $x = \{x_{w_1}, \dots, x_{w_{|s_j|}}\}$. See Figure 1. We defined this regulatory process such that it is derived from three information sources. We have the constant vector φ_S derived from the whole corpus S and from a subset of sentences in it sharing each word w_i . Also we have the vector $\varphi_{w_i}^{(s_j)}$ which is derived from the subset of words $w_i \in V$ such that they take place in $s_j(\cdot, \cdot)$. This formulation proposes that word embeddings x_{w_i} composing a sentence representation have different importances or contributions drawn by some abstract random process of communication, the language.

As an example, we show a sketch of the sentence “*The dog barks*”. Our model (12) allows to see how some $s(\cdot, \cdot)$ would look like geometrically (Figure 2). The sentence sketch could have weights like $\varphi_{The} = 0.075$, $\varphi_{dog} = 0.53$, $\varphi_{barks} = 0.37$. Thus, the example $s[(x_{The}, x_{dog}, x_{barks}), \varphi]$ can be represented as

$$\begin{aligned} s(x, \varphi) &= s[(x_{The}, x_{dog}, x_{barks}), \varphi] = \sum_{w_i \in s} \langle \varphi_{w_i}^{(s)}, \varphi_S \rangle x_{w_i} \\ &= 0.075x_{The} + 0.53x_{dog} + 0.37x_{barks} \end{aligned}$$

The modularity of our model allows for incorporating other information sources through its weights as required, e.g. syntax and structure (Těšitelová, 1992; De Marcken, 1999; Ferrero et al., 2017).

Time and memory cost. Suppose we have an NLP system that relies on sentence representations. Once we have pretrained word embeddings and entropy weights (e.g. TF-IDF), the sentence representations do not need to be explicitly

inferred before they are needed. The representations can be computed *online* as the application requires them. this is possible due to *i*) the modularity of our model and *ii*) the low-cost operations needed to embed a sentence representation. Given a sentence of n words required by an NLP system, WISSE needs d scalar multiplications to weight each word embedding. WISSE also needs the summation of the n weighted word embeddings. Therefore the time needed for each sentence representation is $T_{mul}(dn) + T_{sum}(nd) = T_s(2n + 2d)$. Given that the average sentence length is very short ($n \in \{8, 12\}$ words), we see that $n \ll d$, so $T_s(2n + 2d) \rightarrow \mathcal{O}(kd)$, which is a linear bound to the dimension of the word embeddings. Regarding the memory requirements, the elements of the model can be indexed (e.g. by a database). In this case, for each sentence we only need to load n embeddings of d dimensions and the corresponding n weights, i.e. $n + nd = n(1 + d)$.

7. Design of experiments

In this section we present the multiple aspects considered in designing our experiments. The datasets we used are described, followed by a description of the set of hyperparameters that we used for model selection.

7.1. Datasets

We accomplished comparisons between our method and state-of-the-art methods on the basis of the SICK and the SemEval (2016) datasets (Bentivogli et al., 2016; Agirre et al., 2016). Evaluations of some of these methods have been performed recently on the SemEval (2016) STS competitions (King et al., 2016; Cer et al., 2017).

The SICK and the SemEval (2016) datasets collect sentence pairs from multiple sources. These datasets contain texts written by Internet users, so they are linguistically varied and raise a number of interesting challenges for research in unsupervised sentence representation methods. An important part of the datasets comes from repositories like forums and news appearing daily. The datasets also consider digital dictionaries and knowledge bases (e.g. Wikipedia, WordNet, FrameNet, OntoNet and others). In fact, the SICK dataset is a combination of carefully selected sentences from past years' SemEval STS datasets (2012-2014).

In all cases, the pairs of sentences are associated with an averaged similarity gold standard score manually annotated. The score ranges from 0.0, for unrelated pairs, to 5.0, for equivalent or literally equal pairs. The details about the collection protocol of the dataset are described elsewhere (Agirre et al., 2016; Bentivogli et al., 2016). We evaluated WISSE by taking advantage of this gold standard. A general summary of the datasets is shown below.

The Answer-Answer dataset (SemEval-2016). This dataset contains 1572 pairs of answers which were extracted from the Stack Exchange Data Dump. The pairs of answers correspond to technical Stack Exchange forums, such as academia, cooking, coffee, DIY, and so on.

The Headlines dataset (SemEval-2016). This dataset contains 1498 pair of news headlines collected from the Europe Media Monitor (EMM).

The Postediting dataset (SemEval-2016). This dataset contains 3287 sentences from manually corrected machine translations of English-Spanish-French news. The translations were made using the Moses machine translation system. After that, the translations were manually post-edited so as to be corrected.

The Plagiarism dataset (SemEval-2016). This dataset contains 1271 short answers to computer science questions that exhibit varying degrees of plagiarism with respect to Wikipedia articles.

The Question-Question dataset (SemEval-2016). This dataset contains 1555 pairs of questions from the Stack Exchange Data Dump. The questions were extracted from technical Stack Exchange sites, such as academia, cooking, coffee, DIY, and so on.

The OnWN dataset (SemEval-2013). This dataset contains 561 pairs of sense definitions of terms from WordNet and OntoNotes.

The FNWN dataset (SemEval-2013). This dataset contains 189 pairs of sense definitions of terms from WordNet and FrameNet.

The SICK dataset (SemEval-2014). This dataset contains 4906 pairs of sentences selected from the STS Benchmark. This is a selection of the English datasets used in previous STS tasks (2012-2014). The selection is centered at the difficulty level in STS prediction.

The Wikipedia dataset (general-purpose and unlabeled data). This dataset contains 5.11 million word types and it was downloaded from the Wikipedia dump (2012).

7.2. Experimental Setup

The initial preparation for our experiments was to train the word embedding algorithms. For both W2V⁷ and FastText⁸ word embedding models we used dimensions specified in Table 2, as well as their usual training hyperparameters. The most important of them are the context window length (equal to 10) and the minimum word frequency (equal to 2). These word embedding models were trained on the Wikipedia dataset. For Glove we used 100, 200 and 300 dimensions pretrained embeddings available at the authors' website⁹. For Dep2Vec we used the 300-dimensional pretrained embeddings available at the authors' website¹⁰.

⁷<https://code.google.com/archive/p/word2vec>

⁸<https://github.com/facebookresearch/fastText>

⁹<https://nlp.stanford.edu/projects/glove>

¹⁰<http://u.cs.biu.ac.il/~yogo/data/syntemb/deps.words.bz2>

With respect to information weights, we trained a TF-IDF model¹¹ on the Wikipedia for global training and on each STS dataset for local training. We used default settings, except that we explicitly specified in Table 2.

So as to obtain a reliable idea of the possibilities of WISSE, we performed a number of experiments using datasets listed in Section 7.1. To this end, we selected the hyperparameters of our model (WISSE). Their effects on the model are summarized in Table 2. Combinations of such hyperparameters result in different versions of our model. Thus the idea is to select the best versions.

Overall, both for SICK and SemEval (2016) datasets we first tuned the hyperparameters of our model (e.g. embedding dimension, TF-IDF weighting, and so on). State-of-the-art sentence representation methods have already been evaluated in previous work over these datasets (Section 7.1). For each particular dataset we compared the Pearson coefficient obtained by state-of-the-art methods with that obtained by best versions of WISSE. This allowed us to compare the behavior of these methods, through varied scenarios.

Table 2: Hyperparameters tuned to study WISSE’s performance.

Name	Description	Values
Dataset	The dataset for evaluation	Plagiarism, Answer-Answer, Headlines, Postediting, Question-Question, FNWN, OnWN and SICK (see Section 7.1)
Embedding	Word embedding method	W2V, Glove, FastText and Dep2Vec
Size	The embedding dimensions	100d, 200d, 300d, 400d, 500d, 700d, 1000d.
Comb	Combination method to obtain each sentence representation	Summation (sum), average (avg).
Weights	Weighting method giving word embedding coefficients. All weighting schemes were optionally computed with stop words removal. To denote it, the suffix -st was added. In the same way, the suffixes -bin and -log were added to denote whether the TF vector was computed either as binary or as the logarithm of word frequency.	<ul style="list-style-type: none"> • TF-IDF global –with Wikipedia (<code>glob-tfidf</code>), • IDF global (<code>glob-idf</code>), • TF-IDF local –with STS dataset (<code>loc-tfidf</code>), • IDF local (<code>loc-idf</code>), • All weights equal to 1.0 (unweighted).
Distance	The distance to measure similarity between sentence vectors	Cosine, Euclidean, Manhattan.

In addition to the hyperparameters tuned for our model, we decided to test three different similarity functions for computing semantic similarity scores from WISSE sentence representations: the cosine similarity, the Euclidean distance and the Manhattan distance. For these three functions, the result is directly used as a predicted similarity score¹² (Section 3).

¹¹http://scikit-learn.org/stable/modules/feature_extraction.html

¹²As the Pearson’s coefficient is scale-invariant it is not needed to transform the cosine similarities (which are $[-1.0, 1.0]$) or Euclidean and Manhattan distances (which are $[0.0, \infty)$) into the range of similarities of the gold standard (which are $[0.0, 5.0]$).

8. Experimental Results

Our results are divided mainly into two parts. We presented first the results of the experiments for the SICK dataset. At this stage, we tuned hyperparameters of the WISSE model. After that, we presented the comparison between our best result and the best results obtained by the state-of-the-art unsupervised sentence representation models.

As in the case of SICK, we also tuned WISSE hyperparameters for the SemEval (2016) dataset. Accordingly, the best results obtained with WISSE were compared with the results obtained with state-of-the-art methods. Since the SemEval dataset comprises other datasets (Section 7.1), we presented our best results obtained for each of the SemEval STS datasets (2016).

Finally, we present a statistical comparison between our best result and the best results obtained by the state-of-the-art methods on the SemEval STS datasets. This last comparison includes not only unsupervised sentence representation models, but also supervised representation models as well as supervised and unsupervised STS systems.

8.1. Hyperparameter tuning on the SICK dataset

The combination of all parameters of the model gave us a ranking of up to 200 experiments for the SICK dataset. The best 10 hyperparameter combinations for the three similarity functions (cosine, Euclidean and Manhattan) are shown in Table 3.

Table 3: WISSE hyperparameter combination results on the SICK dataset (train+test)

Weights	Comb.	Size	Embedding	Cosine ρ	Euclid. ρ	Manhatt. ρ
glob-tfidf-bin-st	sum	300d	FastText	0.72405	0.64465	0.64387
glob-tfidf	sum	200d	FastText	0.72023	0.64657	0.6469
glob-tfidf-bin	sum	300d	W2V	0.71995	0.66747	0.66751
loc-tfidf-log	sum	350d	FastText	0.71905	0.65583	0.65615
loc-tfidf-bin	sum	400d	FastText	0.71852	0.65236	0.65209
loc-tfidf-st	avg	300d	Glove	0.70397	0.61817	0.61885
glob-tfidf-st	sum	300d	Dep2Vec	0.67925	0.60972	0.61018
loc-tfidf-log	avg	300d	W2V	0.67428	0.6199	0.62075
loc-tfidf-bin-st	avg	300d	W2V	0.66410	0.58308	0.58289
loc-tfidf-log	avg	300d	Dep2Vec	0.64762	0.55620	0.55585

Bold value indicates the best result.

Our model reached its maximum correlation of $\rho = 0.72405$ for the cosine similarity function. The IDF vector was globally learned (**glob**) from the Wikipedia corpus, which was also used for training the word embeddings. The TF vector was derived binary (**bin**) from the words in each sentence pair, i.e. $\varphi_{w_i}^{s_j} = \{f_{ij} > 0 ? 1.0/F : 0.0\}$. Moreover, in this particular experiment, we stripped out stop words (**-st**) and, therefore, they were omitted from the obtained sentence representations. **FastText** embeddings of 300 dimensions (300d

size) resulted in the best ones over `W2V`, `Glove` and `Dep2Vec`. The weighted word embeddings were combined by adding them (`sum`). We also performed experiments using a simple word embedding average (`avg`). It did not give us higher results than the state-of-the-art mean. Notice that for this dataset, no other similarity function was better than the `cosine`. Therefore, this dataset seems to be best characterized by the angles between word embeddings in sentence representations.

8.2. Results for the SICK dataset

The main result presented in this paper is the performance of our model with respect to the state of the art on unsupervised sentence representation methods.

Table 4: Performance of sentence representation models on the SICK dataset

Sentence representation method (reference)	Pearson
Glove+WR (Arora et al., 2017)	0.722
Sent2vec (Pagliardini et al., 2017)	0.720
FastSent (Hill et al., 2016)	0.720
C-PHRASE (Pham et al., 2015)	0.720
CHARAGRAM-PHRASE (Wieting et al., 2016)	0.700
Skip-thoughts (Kiros et al., 2015)	0.600
BoW TF-IDF (Salton et al., 1983)	0.580
SDAE (Hill et al., 2016)	0.460
Doc2Vec (Le and Mikolov, 2014)	0.460
SAE Hill et al. (2016)	0.310
WISSE (glob-tfidf-bin-st,sum,300d,FastText)	0.724

Bold value indicates the best result.

One of the most popular tasks for evaluating unsupervised sentence representations is using them to predict similarity scores of the SICK dataset (Bentivogli et al., 2016). Therefore, we compared the correlation given by our best model and the correlation given by state-of-the-art models (Table 4).

To our knowledge, both Arora et al. (2017) and Pagliardini et al. (2017) are the best unsupervised sentence representation methods at this moment. Additionally we collected results obtained by other unsupervised methods and the Bag-of-Words baseline (which uses binary TF). Based on the Pearson’s correlation coefficient, WISSE performed better than all them ($\rho = 0.724$).

It is important to notice that the difference is not large between WISSE and the following two methods in the ranking (Table 4). There are now 5 sentence embedding methods within such a barrier (Glove+WR, Sent2vec, FastSent, C-PHRASE and WISSE).

8.3. Hyperparameter tuning on the SemEval STS datasets

We obtained correlation coefficients for WISSE on the SemEval STS datasets (2016) (see Table 5). These datasets together present much more variety with

respect to the SICK dataset. Therefore, the results also varied independently of the evaluated model.

Our best result was obtained for the Postediting dataset ($\rho = 0.82161$). The `tfidf` weights were computed locally from the Postediting dataset by using word frequencies as TFs. Stop words were included in the model. The best embeddings for this dataset were Dep2Vec of 300 dimensions. The Euclidean distance better captures the similarity between sentence representations. Notice that there is a relatively large difference between the metrics (Euclidean and Manhattan) and the cosine similarity. Postediting was the unique dataset for which Dep2Vec embeddings led to the best result.

Our second best result was obtained for the Plagiarism dataset ($\rho = 0.80607$). For this dataset the best hyperparameters varied considerably with respect to the best for Postediting. The main differences were observed in word embedding weighting and word embeddings. The `tfidf` weights were computed globally from Wikipedia and by using binary TFs. The best word embeddings were W2V of 300 dimensions.

Our lowest result in SemEval (2016) was for the Answer-Answer dataset ($\rho = 0.6556$). This dataset has been reported to be challenging for most STS competitors (Agirre et al., 2016). Our model needed W2V embeddings of 1000 dimensions. The weights of the model were computed locally and by using the logarithm of the word frequencies in the sentences, i.e. $\varphi_{w_i}^{(s_j)} = \log(f_{ij} + 1)/F$.

Table 5: WISSE hyperparameter combination results on the SemEval STS dataset (2016)

Dataset	Embedding	Size	Weights	Cosine	Euclidean	Manhatt.
Postediting	Dep2Vec	300d	loc-tfidf	0.652880	0.821610	0.819890
Plagiarism	W2V	300d	glob-tfidf-bin	0.775750	0.806070	0.805280
Ques.-Ques.	FastText	300d	glob-tfidf	0.704010	0.683410	0.681400
Headlines	FastText	200d	glob-tfidf	0.676300	0.701020	0.700720
Ans.-Ans.	W2V	1000d	loc-tfidf-log	0.507660	0.655600	0.652110
OnWN	W2V	1000d	loc-tfidf-log-st	0.833070	0.739430	0.738610
FNWN	FastText	200d	glob-tfidf	0.458560	0.350300	0.363520

For all cases the best combination operation was the summation (`sum`).

Bold value indicates the best result.

In most cases, the two metrics (Euclidean and Manhattan) outperformed the cosine similarity and the difference is relatively large. We observed this for the majority of the datasets (Table 5): Postediting, Plagiarism, Headlines and Answer-Answer. For all datasets, the weighted word embeddings were summed in order to obtain the corresponding sentence representation. In fact, all our best results were obtained by summing the weighted embeddings (there were no cases where the word embedding average (`avg`) performed better).

The dimension of the embeddings did not show specific patterns, except for the OnWN and Answer-Answer datasets. Both of them needed logarithmic TF vectors as well as 1000 dimensional word embeddings. In contrast to the other datasets, these datasets allowed incremental dimensionality augmen-

tation (200 – 1000). Nonetheless, once such an augmentation surpassed 500 dimensions the performance almost did not improve. For all other datasets, augmenting dimensions degraded the performance. WISSE performed better for all the datasets when the weights were the TF-IDF vector products (using either binaries, frequencies or their logarithm) instead of only IDF vector components. This observation could be considered independent of the word embedding method and dimension.

As an additional test we evaluated our model on the OnWN, FNWN (SemEval 2013) datasets. These datasets have been reported to be challenging (Agirre et al., 2013).

On the one hand the OnWN dataset was shown to be well represented by the BoW baseline, $\rho = 0.8431$; however for SemEval competing systems it is hard to surpass such a performance. Our best result was comparable ($\rho = 0.82161$) by using W2V embeddings of 1000 dimensions. The IDF vectors were locally learned and the TF vectors were computed from the logarithm of the word frequency in sentences. The stop words were stripped out.

On the other hand, for the FNWN dataset the competing systems were struggling to reach the maximum $\rho = 0.5818$. Our system reached $\rho = 0.45856$ by using FastText word embeddings of 200 dimensions. The embeddings were weighted via globally learned IDFs and frequency TFs. For both datasets, the cosine similarity worked considerably better than the metrics.

8.4. Results for the SemEval STS task

We compared WISSE with the sentence embedding methods from the state of the art on the SemEval STS dataset (2016). For WISSE, we have used the best hyperparameters fitted for each dataset (Section 8.3).

Table 6: WISSE and state-of-the-art sentence embedding models on SemEval STS datasets

Model	Ans.-Ans.	HDL	Plagiarism	Postediting	Ques.-Ques.
Sent2Vec ¹³	0.57739	<i>0.75061</i>	0.80068	<i>0.82857</i>	<i>0.73035</i>
WISSE	0.65560	0.70102	0.80607	0.82161	0.70410
D2V (400d)	0.41123	0.69169	0.60488	0.75547	-0.02245
Skip-toughs	0.23199	0.49643	0.48636	0.17749	0.33446
W2V (300d-average)	0.50311	0.66362	0.72347	0.73935	0.16586
Binary TF-IDF (BoW)	0.41133	0.54073	0.69601	0.82615	0.03844

Bold values indicates the best result for WISSE.

Values in bold italics indicate best state-of-the-art results.

We observed that for two datasets (Answer-Answer and Plagiarism) WISSE performed better than the state of the art methods (Table 6). For the other three datasets (Headlines (HDL), Postediting and Question-Question) the Sent2vec method was better. We observed that, in general, it is difficult to beat the BoW baseline in the Postediting dataset. Sent2Vec surpassed it by a small difference. WISSE was lower, also, by a small difference.

¹³This method was tested by the MayoNLP team (Afzal et al., 2016) .

Although currently D2V (Paragraph vector) and Skip-Thoughts are very popular methods, their performances in STS were even lower than that of the BoW baseline (excepting the Postediting dataset). The performances of such methods were also considerably lower than those of Sent2Vec and WISSE. For instance, while Sent2vec and WISSE reached roughly $\rho > 0.70$ on the Question-Question dataset, D2V fell to $\rho = 0.02245$. A similar situation occurred for the Skip-Thoughts method, which reached $\rho = 0.17749$ on the Postediting dataset while the best methods were $\rho > 0.82$. Finally, the simple average of W2V word embeddings in building sentence representations showed to be a method with very similar performance with respect to results obtained with D2V.

8.5. The SemEval STS task

We performed a comparison between WISSE and the top 10 STS systems in the SemEval STS (2016) competition (Table 7). Although we did not participate in such a competition, it is relevant to observe how WISSE performed with respect to state-of-the-art STS systems. Notice that these systems can be designed in varied forms and for different purposes than WISSE (Section 3). That is, while WISSE aims to unsupervisedly represent sentences in vector spaces, STS systems aim to measure semantic textual similarity (un/supervisedly and by using varied external resources).

Table 7: Ranking for the SemEval STS 2016 task (this is a modified version derived from Agirre et al. (2016))

R	System	ALL	Ans.-Ans.	HDL	Plagiarism	Postediting	Ques.-Ques.
1	Samsung Pol.	0.77807	0.69235	0.82749	0.84138	0.83516	0.68705
2	UWB	0.75731	0.62148	0.81886	0.82355	0.82085	0.70199
3	MayoNLPTeam	0.75607	0.61426	0.77263	0.805	0.8484	0.74705
4	Samsung Pol.	0.75468	0.69235	0.82749	0.81288	0.83516	0.58567
5	NaCTeM	0.74865	0.60237	0.8046	0.81478	0.82858	0.69367
6	ECNU	0.75079	0.56979	0.81214	0.82503	0.82342	0.73116
7	UMD-TTIC-UW	0.74201	0.66074	0.79457	0.81541	0.80939	0.61872
9	SimiHawk	0.73774	0.59237	0.81419	0.80566	0.82179	0.65048
8	Sent2Vec	0.73836	0.57739	0.75061	0.80068	0.82857	0.73035
10	WISSE	0.73768	0.655600	0.70102	0.80607	0.82065	0.70410
23	UWB	0.72622	0.64442	0.79352	0.82742	0.81209	0.53383
-	D2V (400d)	0.50206	0.41123	0.69169	0.60488	0.75547	-0.02245
-	Skip-toughs	0.27148	0.23199	0.49643	0.48636	0.17749	0.33446
-	W2V (300d-avg)	0.56007	0.50311	0.66362	0.72347	0.73935	0.16586
110	STS (BoW)	0.51334	0.41133	0.54073	0.69601	0.82615	0.03844

We used our best hyperparameters for each dataset (Section 8.3). The general score for the competition is a weighted average of the correlations from all datasets (column ALL). It is encouraging that WISSE is ranked at 10th place among 113 systems. This comparison was performed putting the characteristics of all those systems aside. This is an important consideration because most of them are supervised and include several lexical resources like WordNet, FrameNet and dictionaries. Thus, the unsupervised sentence representation models are separated at the middle of Table 7 in order to keep such a distinction.

To know whether our model is actually competitive, we analyzed the correlation coefficients given by participant systems (Figure 3) for the overall performances (ALL) and the performance for each dataset (Table 7). The first

observation is that the WISSE model (green diamond) outperformed the mean for almost all datasets. For the Headlines dataset, WISSE did not surpass the mean, but it stayed within the inter quartile range (IQR). Neither the Plagiarism nor the Postediting datasets, do not present high difficulties for the state-of-the-art systems. This can be a result of the low variability and the high overall mean correlation reached by most systems for this dataset. In the same vein, the BoW baseline (red circle) overpasses the mean correlation for the Postediting dataset. For these datasets, our system remains within the IQR, but not too far away from the maximum correlation.

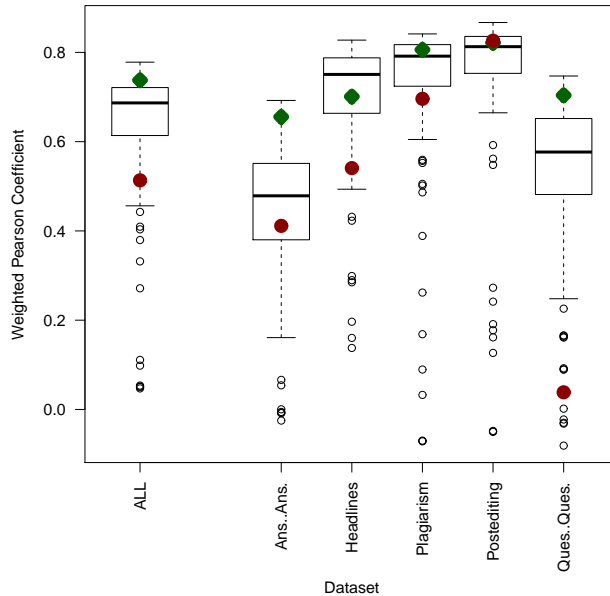


Figure 3: Box plot of the statistical position of the WISSE model within the SemEval STS 2016 state of the art. The vertical axis is the correlation coefficient given by participant systems for each dataset of the STS competition. The horizontal axis shows the overall performances (ALL) and the performance for each dataset. (WISSE = green diamonds; BoW baseline = red circles).

Statistically, our best results were reached on both the Answer-Answer and the Question-Question datasets. Our model was allocated beyond the first quartile, near to the maximum overall performances. That is, WISSE succeed in dealing with the uncertainty of the mean correlations reached by the 95% of the systems (IQR) on such difficult datasets. For the Question-Question dataset the baseline is at the lower outlier zone. The Answer-Answer dataset is one of the most difficult datasets to predict for the state of the art methods. Nonetheless, even the BoW baseline reaches the IQR for this dataset. In the general state-of-the-art correlation (ALL), WISSE is located barely beyond the first quartile;

in fact within the upper whisker of the box plot.

9. Discussion

The differences in the behavior of our sentence representation method with respect to the textual properties of each dataset are meaningful. As confirmed by our results, different genres of text give different results.

The fact that WISSE outperformed the state of the art on the SICK dataset is interesting. This is because such a dataset constitutes a careful selection of sentence pairs coming from varied STS datasets (2012-2014). This selection included cross-level tasks (Jurgens et al., 2016), which are the most difficult tasks in existent STS benchmarks. In fact, we observed that it is very difficult to surpass the 0.7200 barrier and we think that it is mainly due to cross-level samples. These kind of samples are mostly used for textual entailment tasks, rather than for learning sentence representations. Nonetheless, supervised methods have been successful in learning generalization of both kind of tasks simultaneously (Yin et al., 2016). There are now 5 sentence embedding methods within such a barrier (Glove+WR, Sent2vec, FastSent, C-PHRASE and WISSE). All these methods are completely based on neural networks. Our method (WISSE) integrates neural networks and Shannon’s entropy, the former for learning word embeddings, and the second for weighting word embeddings.

In addition, the experiments on diverse text properties presented by the SemEval (2016) dataset showed that WISSE offers the possibility to be prepared for varied scenarios. Such preparation consists of using proper hyperparameters for text properties, which is possible due to the modularity of our model.

Shannon’s entropy of words in the corpus can be used for weighting the contribution of each word embedding. Therefore, these two elements (words and their information contributions) are identifiable at the time that our method builds sentence representations. This is a partial¹⁴ identifiability allowing for future studies on statistical properties of sentence meaning.

A surprising aspect of our approach is that the entropy is indeed a real number (a scalar). This fact differs considerably from purely neural network-based approaches: a whole matrix should be learned in order to capture interactions between word embeddings comprising a sentence representation. Thus our experiments were performed at relatively low computational cost.

For each dataset we wanted to know whether our model is actually competitive with respect to the overall STS state of the art. We observed in Section 3 that there is a wide variety of STS systems. Several of them are based on knowledge and language resources as well as on manually annotated similarity gold standards for supervised learning. As we reported in Section 8.5, these advantages do not always represent for a significant advantage over unsupervised sentence representation methods. In this sense, in most cases presented

¹⁴Partial identifiability is referred to as the fact that we are using word embeddings that do not come necessarily from an identifiable model.

in this paper, WISSE succeed in dealing with the uncertainty of the mean correlations reached by 95% of all state-of-the-art systems. This holds even when the variability of the overall performances is high.

We think that the aforementioned facts indicate that WISSE is a competitive method for representing meaning in sentences.

There are datasets where most methods in the state of the art (including WISSE) reached low performances, i.e. FNWN and Answer-Answer. For the case of WISSE, we think this was because of the magnitude that the sentence representations attained due to the contributions of weighted word embeddings. It subsumes the word embeddings associated with the words composing a text snippet. The larger the number of words contained within, the larger the magnitude reached by the resulting representation¹⁵. The FNWN and Answer-Answer datasets constitute cross-level STS tasks (Jurgens et al., 2016), e.g. the comparison between a phrase and a paragraph. The length imbalance makes WISSE producing significantly different vectors to be compared. For instance, even when the meanings represented are very similar, length imbalance can cause meaningful geometric differences. This fact was confirmed as the cosine similarity was the best similarity function for the datasets in question (Section 8.3). In these cases we think that there were more useful similarity cues encoded in the angles of vectors than into their lengths (which are actually omitted by the cosine similarity). Conversely, for balanced length between compared snippets the metrics led to better results because such metrics exploited the fact that sentence meaning was encoded into the sentence representation magnitudes (and implicitly into the angles).

10. Possible improvements

Blinded information sources. Our current experiments only considered a basic version of our model using three sources of information provided by Shannon’s entropy (TF-IDF). Nonetheless, as we can see in eqs. (5) and (6), the usual IDF weights make naïve assumptions in terms of the probability measures of words and sentences. We think these assumptions might circumvent important things. The most important is the actual probability measures of the stochastic processes underlying the different levels and hierarchies of contexts. Clearly, uniform distributions are not the way sentences are drawn from the communication process. In this sense, we think that more accurate empirical estimations should be performed in order to uncover more accurate models of stochastic processes underlying language (Vapnik, 1998). These estimations will provide us more granular information sources, better weighting estimates and probably better sentence representations (of course, it will depend on the requirements of the addressed NLP task).

¹⁵This affirmation is relative. In ℓ_p spaces, with $1 < p < \infty$, it holds that $\|x\|_p = (\sum_{i=1}^{\infty} |x_i|^p)^{\frac{1}{p}} < \infty$.

Consider the case of *word order universals* (Derbyshire, 1977), which is an information source we have left blinded for the current version of our model. A speaker emitting a sentence like (3) generally emphasizes the agent (*Paula*) performing the action (*to strike*). Conversely, in sentences like (4), the speaker emphasizes the experiencer (*Nacho*). This is not a general rule, but it surely provides an information spectrum related to informativeness of noun order.

(3) *Paula struck Nacho.*

(4) *Nacho was beaten by Paula.*

Notice that self similarity and hierarchy of these patterns can be found also in more complex constructions, e.g., *Paula struck Nacho because his offense was inadmissible*. In this example, the noun phrase *his offense* in the subordinate clause is also emphasized by the speaker (with respect to the adjective *inadmissible*). Nonetheless, for the whole sentence this emphasis is hierarchically less important than that put onto the agent in the main clause (*Paula*).

Advantages and disadvantages. At the moment, WISSE does not capture properly specific information patterns, like negation adverbs. For instance, similar representations of the sentences (1) and (2) are built:

(1) *Computers will not condemn humanity.*

(2) *Computers will condemn humanity.*

This holds because the same adverb form (*not*) is used many times when the speaker needs to negate any given fact. Thus, from the point of view of the information embedding that WISSE uses at this moment, this word is almost not informative. We show in Table 8 different weights for the word *not* within two randomly picked sentences from our corpus:

Table 8: TF-IDF weights for the adverb form *not* within two sentences

Sentence	TF-IDF weights
<i>The man jumping is not wearing a shirt.</i>	“jumping” (0.6537), “wearing” (0.4753), “shirt” (0.5679), “not” (0.1894)
<i>A girl is close to a boy whose face is not shown.</i>	“girl” (0.4258), “close” (0.3415), “boy” (0.4339), “face” (0.3833), “not” (0.1959)

In fact, negation is an open issue in STS research. Depending on the way we consider the TF (either binary or logarithmic or frequency), the weights will change. Nonetheless their entropy keeps relatively constant.

Notice that we are clearly identifying specific linguistic items participating in our model. This is an advantage of WISSE over most sentence representation methods. It is possible to manipulate the contribution of specific linguistic items within sentences. For instance, we could be interested in detecting similarity of sentences with respect to a specific entity (a noun phrase). In this case, it is

sufficient to reinforce the weights associated with the word embeddings representing such an entity (Badarınza et al., 2017). This identifiability is not easy to attain by other state-of-the-art purely neural sentence embedding methods.

11. Conclusions

We evaluated our approach on well-known STS benchmarks. Our method outperformed the usual baselines of sentence representation (BoW and word embedding average). Furthermore, the evaluation on the SICK dataset showed that our method outperformed the state of the art methods by reaching a correlation of $\rho = 0.724$. This is encouraging because this dataset is overall a bit difficult for unsupervised sentence representation methods. In fact, the barrier of 0.72 is difficult to overcome.

WISSE also outperformed state-of-the-art methods in a couple of datasets (Answer-Answer and Plagiarism) of the well-known SemEval (2016) STS benchmark. Particularly for the Answer-Answer dataset our method surpassed the state of the art methods by about 10%, although this is one of the most challenging datasets. For the remaining datasets, our method’s performance was close to the state of the art methods (within 2.7%).

We also compared WISSE against a wide variety of methods that participated in the SemEval STS competition (2016). Although such methods were not restricted to unsupervised or external knowledge-free methods, our model surpassed the mean performance. Furthermore, in most cases WISSE surpassed 95% of the competitors and it was close (within 4%) to the best STS system (which uses supervised learning, i.e. support vector regression and deep learning, as well as a variety of external resources such as WordNet, Wikipedia and named entity recognizers).

Our experiments confirmed our hypothesis that it is possible to well represent sentences by using the link between the contexts learned by word embeddings and the entropy of embedded words. We exploited the mentioned link to learn without supervision the weights of a series of word embeddings representing a sentence. Interestingly, such weights are simple scalars that allowed our model to reach state-of-the-art performance in difficult STS tasks at low computational cost.

The modularity of our model gives the possibility of configuring our model according to the properties of text, which allowed us to obtain the best performance in important STS tasks. It is also interesting that such performances were obtained by using simple similarity functions (cosine, Euclidean, Manhattan).

Finally, the low computational cost of our method makes it especially useful when only unlabeled text is available for learning sentence representations. Furthermore, our model is suitable for low-resource applications and for leveraging real-time and online analysis and applications of sentence representations.

Acknowledgements

We thank to: CONACyT (Grant 386128), *Laboratorio de cómputo de alto rendimiento* (IIMAS–UNAM), Alfredo José Hernández Álvarez (CCG–UNAM), The Computational Genomics Research Program (PGC/CCG–UNAM) and The Postgraduate Program in Computer Science and Engineering–UNAM. Also thanks to Sarah Elizabeth Champion (Qatar University) for helping others in an unselfish way.

References

- Afzal N, Wang Y, Liu H. Mayonlp at semeval-2016 task 1: Semantic textual similarity based on lexical semantic net and deep learning semantic model. In: SemEval NAACL-HLT. 2016. p. 674–9.
- Agirre E, Banea C, Cer D, Diab M, Gonzalez-Agirre A, Mihalcea R, Rigau G, Wiebe J. Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation. In: Proceedings of SemEval. 2016. p. 497–511.
- Agirre E, Cer D, Diab M, Gonzalez-Agirre A, Guo W. * 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. volume 1; 2013. p. 32–43.
- Agirre E, Diab M, Cer D, Gonzalez-Agirre A. Semeval-2012 task 6: A pilot on semantic textual similarity. 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. Association for Computational Linguistics; 2012. p. 385–93.
- Aizawa A. An information-theoretic perspective of tf–idf measures. *Information Processing & Management* 2003;39(1):45–65.
- Arora S, Liang Y, Ma T. A simple but tough-to-beat baseline for sentence embeddings. *International Conference on Learning Representations (ICLR)* 2017;.
- Arroyo-Fernández I. Learning kernels for semantic clustering: A deep approach. In: NAACL-HLT 2015 Student Research Workshop (SRW). 2015. p. 79–87.
- Arroyo-Fernández I, Meza Ruiz IV. LIPN-IIMAS at SemEval-2017 task 1: Subword embeddings, attention recurrent neural networks and cross word alignment for semantic textual similarity. In: Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017). Vancouver, Canada: Association for Computational Linguistics (ACL); 2017. p. 199–203. URL: <http://www.aclweb.org/anthology/S17-2031>.

- Arroyo-Fernández I, Torres-Moreno JM, Sierra G, Cabrera-Diego LA. Automatic text summarization by non-topic relevance estimation. In: Proceedings of the 8th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management - Volume 1: KDIR., 2016. p. 89–100. doi:10.5220/0006053400890100.
- Badarinza I, Sterca AI, Ionescu M. Syntactic indexes for text retrieval. *Information Technology in Industry* 2017;5:24–8.
- Baroni M, Dinu G, Kruszewski G. Don't count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In: *ACL (1)*. 2014. p. 238–47.
- Baroni M, Lenci A. Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics* 2010;36(4):673–721.
- Bengio Y, Ducharme R, Vincent P, Jauvin C. A neural probabilistic language model. *Journal of Machine Learning Research* 2003;3(Feb):1137–55.
- Bentivogli L, Bernardi R, Marelli M, Menini S, Baroni M, Zamparelli R. Sick through the semeval glasses. lesson learned from the evaluation of compositional distributional semantic models on full sentences through semantic relatedness and textual entailment. *Language Resources and Evaluation* 2016;50(1):95–124.
- Bojanowski P, Grave E, Joulin A, Mikolov T. Enriching word vectors with subword information. *arXiv preprint arXiv:160704606* 2016;.
- Brokos GI, Malakasiotis P, Androutsopoulos I. Using centroids of word embeddings and word movers distance for biomedical document retrieval in question answering. *ACL 2016* 2016;:114.
- Brychcin T, Svoboda L. Uwb at semeval-2016 task 1: Semantic textual similarity using lexical, syntactic, and semantic information. In: *Proceedings of SemEval*. *ACL 2016*; 2016. p. 588–94.
- Cer D, Diab M, Agirre E, Lopez-Gazpio I, Specia L. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In: *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. Vancouver, Canada: Association for Computational Linguistics; 2017. p. 1–14. URL: <http://www.aclweb.org/anthology/S17-2001>.
- Charniak E. *Statistical language learning*. MIT Press, 1996.
- Chen T, Xu R, He Y, Wang X. Improving sentiment analysis via sentence type classification using bilstm-crf and cnn. *Expert Systems with Applications* 2017;72:221–30.
- Collobert R, Weston J, Bottou L, Karlen M, Kavukcuoglu K, Kuksa P. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 2011;12(Aug):2493–537.

- De Boom C, Van Canneyt S, Demeester T, Dhoedt B. Representation learning for very short texts using weighted word embedding aggregation. *Pattern Recognition Letters* 2016;80:150–6.
- De Marcken C. On the unsupervised induction of phrase-structure grammars. In: *Natural Language Processing Using Very Large Corpora*. Springer; 1999. p. 191–208.
- Derbyshire DC. Word order universals and the existence of ovs languages. *Linguistic Inquiry* 1977;8(3):590–9. URL: <http://www.jstor.org/stable/4178003>.
- Elman JL. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning* 1991;7(2-3):195–225.
- Er MJ, Zhang Y, Wang N, Pratama M. Attention pooling-based convolutional neural network for sentence modelling. *Information Sciences* 2016;373:388–403.
- Ferrero J, Besacier L, Schwab D, Agnès F. Compilig at semeval-2017 task 1: Cross-language plagiarism detection methods for semantic textual similarity. In: *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. Vancouver, Canada: Association for Computational Linguistics; 2017. p. 100–5. URL: <http://www.aclweb.org/anthology/S17-2012>.
- Firth JR. *Papers in linguistics 1934-1951*. London: Oxford University Press, 1957.
- Fourier J. *Theorie analytique de la chaleur*. Chez Firmin Didot, père et fils, 1822.
- Han L, Kashyap A, Finin T, Mayfield J, Weese J. Umbc ebiquity-core: Semantic textual similarity systems. Atlanta, GA, USA 2013;44.
- Harris ZS. *Mathematical Structures of Language*. New York, NY, USA: Wiley, 1968.
- Hatzivassiloglou V, Klavans JL, Eskin E. Detecting text similarity over short passages: Exploring linguistic feature combinations via machine learning. In: *Proceedings of the 1999 joint sigdat conference on empirical methods in natural language processing and very large corpora*. 1999. p. 203–12.
- Hill F, Cho K, Korhonen A. Learning distributed representations of sentences from unlabelled data. In: *Proceedings of NAACL-HLT*. 2016. p. 1367–77.
- Hinton G, McClelland J, Rumelhart D. Distributed representations. In: *Parallel distributed processing: explorations in the microstructure of cognition*, vol. 1. MIT Press; 1986. p. 77–109.
- Ji Y, Eisenstein J. Discriminative improvements to distributional sentence similarity. In: *EMNLP*. 2013. p. 891–6.

- Jurgens D, Pilehvar MT, Navigli R. Cross level semantic similarity: an evaluation framework for universal measures of similarity. *Language Resources and Evaluation* 2016;50(1):5–33.
- Kalchbrenner N, Grefenstette E, Blunsom P. A convolutional neural network for modelling sentences. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics* 2014;URL: <http://goo.gl/EsQCuC>.
- Kenter T, de Rijke M. Short text similarity with word embeddings. In: *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. ACM; 2015. p. 1411–20.
- King M, Gharbieh W, Park S, Cook P. Unbnlp at semeval-2016 task 1: Semantic textual similarity: A unified framework for semantic processing and evaluation. *Proceedings of SemEval 2016*;:732–5.
- Kintsch W, Mangalath P. The construction of meaning. *Topics in Cognitive Science* 2011;3(2):346–70.
- Kiros R, Zhu Y, Salakhutdinov RR, Zemel R, Urtasun R, Torralba A, Fidler S. Skip-thought vectors. In: *Advances in neural information processing systems*. 2015. p. 3294–302.
- Landauer TK, Foltz PW, Laham D. An introduction to latent semantic analysis. *Discourse Processes* 1998;25(2-3):259–84.
- Le Q, Mikolov T. Distributed representations of sentences and documents. In: *31st International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June. 2014*. p. 1188–96. URL: <http://jmlr.org/proceedings/papers/v32/le14.html>.
- Levy O, Goldberg Y. Dependency-based word embeddings. In: *ACL (2)*. 2014. p. 302–8.
- Manning CD, Raghavan P, Schütze H. *An Introduction to Information Retrieval*. Cambridge, United Kingdom: Cambridge University Press, 2009.
- Martin DI, Berry MW. Mathematical foundations behind latent semantic analysis. *Handbook of latent semantic analysis* 2007;:35–56.
- Meza-Ruiz I, Riedel S. Jointly identifying predicates, arguments and senses using markov logic. In: *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics; 2009. p. 155–63.
- Mihalcea R, Corley C, Strapparava C, et al. Corpus-based and knowledge-based measures of text semantic similarity. In: *AAAI*. volume 6; 2006. p. 775–80.
- Mikolov T, Chen K, Corrado G, Dean J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:13013781* 2013a;.

- Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J. Distributed representations of words and phrases and their compositionality. In: *Advances in Neural Information Processing Systems*. 2013b. p. 3111–9.
- Mitchell J, Lapata M. Composition in distributional models of semantics. *Cognitive Science* 2010;34(34):1388–429. Cognitive Science Society, ISSN: 1551-6709.
- Mitra M, Buckley C, Singhal A, Cardie C. An analysis of statistical and syntactic phrases. In: *Computer-Assisted Information Searching on Internet*. Le Centre de Hautes Etudes Internationales d’Informatique Documentaire; 1997. p. 200–14.
- Onan A, Korukoğlu S, Bulut H. A hybrid ensemble pruning approach based on consensus clustering and multi-objective evolutionary algorithm for sentiment classification. *Information Processing & Management* 2017;53(4):814–33.
- Osteyee DB, Good IJ. Expected mutual information. In: *Information, Weight of Evidence, the Singularity between Probability Measures and Signal Detection*. Berlin, Germany: Springer; 1974. p. 26–38. doi:10.1007/BFb0064132.
- Pagliardini M, Gupta P, Jaggi M. Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features. arXiv 2017;arXiv:1703.02507.
- Pennington J, Socher R, Manning CD. Glove: Global vectors for word representation. In: *Empirical Methods in Natural Language Processing (EMNLP)*. 2014. p. 1532–43. URL: <http://www.aclweb.org/anthology/D14-1162>.
- Pereira F. Formal grammar and information theory: together again? *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 2000;358(1769):1239–53.
- Pham NT, Kruszewski G, Lazaridou A, Baroni M. Jointly optimizing word representations for lexical and sentential tasks with the c-phrase model. In: *ACL (1)*. 2015. p. 971–81.
- Pilehvar MT, Navigli R. From senses to texts: An all-in-one graph-based approach for measuring semantic similarity. *Artificial Intelligence* 2015;228:95 – 128.
- Robertson S. Understanding inverse document frequency: on theoretical arguments for idf. *Journal of Documentation* 2004;60(5):503–20.
- Rong X. word2vec parameter learning explained. *CoRR* 2014;abs/1411.2738. URL: <http://arxiv.org/abs/1411.2738>.
- Rychalska B, Pakulska K, Chodorowska K, Walczak W, Andruszkiewicz P. Samsung Poland NLP Team at SemEval-2016 Task 1: Necessity for diversity; combining recursive autoencoders, Wordnet and ensemble methods to measure semantic similarity. In: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval 2016)*, San Diego, CA, USA. 2016. p. 602–8.

- Salton G, Buckley C. Term-weighting approaches in automatic text retrieval. *Information Processing & Management* 1988;24(5):513–23.
- Salton G, Fox EA, Wu H. Extended Boolean information retrieval. *Communications of the ACM* 1983;26(11):1022–36.
- Schölkopf B, Simard P, Smola A, Vapnik V. Prior knowledge in support vector kernels. In: *Proceedings of the 10th International Conference on Neural Information Processing Systems*. Cambridge, MA: MIT Press; 1997. p. 640–6.
- Shannon CE. Communication theory of secrecy systems*. *Bell system technical journal* 1949;28(4):656–715.
- Spärk Jones K. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation* 1972;28(1):11–21.
- Sultan MA, Bethard S, Sumner T. Back to basics for monolingual alignment: Exploiting word similarity and contextual evidence. *Transactions of the Association for Computational Linguistics* 2014;2:219–30.
- Těšitělová M. Quantitative linguistics. *Linguistics and literary studies in Eastern Europe*. Academia Publishing House of the Czechoslovak Academy of Sciences, 1992.
- Tian R, Okazaki N, Inui K. The mechanism of additive composition. *Machine Learning* 2017;106(7):1083–130.
- Vapnik VN. *Statistical Learning Theory*. New York, NY: Wiley, 1998.
- Wieting J, Bansal M, Gimpel K, Livescu K. Charagram: Embedding words and sentences via character n-grams. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics; 2016. p. 1504–15. URL: <https://aclweb.org/anthology/D16-1157>.
- Yazdani M, Popescu-Belis A. Computing text semantic relatedness using the contents and links of a hypertext encyclopedia. *Artificial Intelligence* 2013;194:176–202.
- Yin W, Schütze H. Discriminative phrase embedding for paraphrase identification. In: *Proceedings of HLT-NAACL*. 2015. p. 1368–73.
- Yin W, Schütze H, Xiang B, Zhou B. Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *Transactions of the Association for Computational Linguistics* 2016;4:259–72.
- Yu J, Xie L, Xiao X, Chng ES. Learning distributed sentence representations for story segmentation. *Signal Processing* 2017;.

Zhang Z, Ge SS, He H. Mutual-reinforcement document summarization using embedded graph based sentence clustering for storytelling. *Information Processing & Management* 2012;48(4):767–78.

Zheng G, Callan J. Learning to reweight terms with distributed representations. In: *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM; 2015. p. 575–84.