Word Embeddings as Tuples of Feature Probabilities

Siddharth Bhat[1], **Alok Debnath**[2], **Souvik Bannerjee**[2] and **Manish Shrivastava**[2]

[1]Center for Security, Theory & Algorithmic Research (C-STAR)

[2]Language Technologies Research Center (LTRC)

Kohli Center on Intelligent Systems

International Institute of Information Technology, Hyderabad

{siddharth.bhat, alok.debnath, souvik.bannerjee}@research.iiit.ac.in

m.shrivastava@iiit.ac.in

Abstract

In this paper, we provide an alternate perspective on word representations, by reinterpreting the dimensions of the vector space of a word embedding as a collection of features. In this reinterpretation, every component of the word vector is normalized against all the word vectors in the vocabulary. This idea now allows us to view each vector as an n-tuple (akin to a fuzzy set), where n is the dimensionality of the word representation and each element represents the probability of the word possessing a feature. Indeed, this representation enables the use fuzzy set theoretic operations, such as union, intersection and difference. Unlike previous attempts, we show that this representation of words provides a notion of similarity which is inherently asymmetric and hence closer to human similarity judgements. We compare the performance of this representation with various benchmarks, and explore some of the unique properties including function word detection, detection of polysemous words, and some insight into the interpretability provided by set theoretic operations.

1 Introduction

Word embedding is one of the most crucial facets of Natural Language Processing (NLP) research. Most non-contextualized word representations aim to provide a distributional view of lexical semantics, known popularly by the adage *"a word is known by the company it keeps"* (Firth, 1957). Popular implementations of word embeddings such as word2vec (Mikolov et al., 2013a) and GloVe (Pennington et al., 2014) aim to represent words as embeddings in a vector space. These embeddings are trained to be oriented such that vectors with higher similarities have higher dot products when normalized. Some of the most common methods of intrinsic evaluation of word embeddings include similarity, analogy and compositionality. While similarity is computed using the notion of dot product, analogy and compositionality use vector addition.

However, distributional representations of words over vector spaces have an inherent lack of interpretablity (Goldberg and Levy, 2014). Furthermore, due to the symmetric nature of the vector space operations for similarity and analogy, which are far from human similarity judgements (Tversky, 1977). Other word representations tried to provide asymmetric notions of similarity in a non-contextualized setting, including Gaussian embeddings (Vilnis and McCallum, 2014) and word similarity by dependency (Gawron, 2014). However, these models could not account for the inherent compositionality of word embeddings (Mikolov et al., 2013b). Moreover, while work has been done on providing entailment for vector space models by entirely reinterpreting word2vec as an entailment based semantic model (Henderson and Popa, 2016), it requires an external notion of compositionality. Finally, word2vec and GloVe, as such, are meaning conflation deficient, meaning that a single word with all its possible meanings is represented by a single vector (Camacho-Collados and Pilehvar, 2018). Sense representation models in non-contextualized representations such as multi-sense skip gram, by performing joint clustering for local word neighbourhood. However, these sense representations are conditioned on non-disambiguated senses in the context and require additional conditioning on the intended senses (Li and Jurafsky, 2015).

In this paper, we aim to answer the question: *Can a single word representation mechanism account for lexical similarity and analogy, compositionality, lexical entailment **and** be used to detect and resolve polysemy?* We find that by performing column-wise normalization of word vectors trained using the word2vec skip-gram negative sampling

regime, we can indeed represent all the above characteristics in a single representation. We interpret a column wise normalized word representation. We now treat these representations as fuzzy sets and can therefore use fuzzy set theoretic operations such as union, intersection, difference, etc. while also being able to succinctly use asymmetric notions of similarity such as K-L divergence and cross entropy. Finally, we show that this representation can highlight syntactic features such as function words, use their properties to detect polysemy, and resolve it qualitatively using the inherent compositionality of this representation.

In order to make these experiments and their results observable in general, we have provided the code which can be used to run these operations. The code can be found at `https://github.com/AlokDebnath/fuzzy_embeddings`. The code also has a working command line interface where users can perform qualitative assessments on the set theoretic operations, similarity, analogy and compositionality which are discussed in the paper.

2 Related Work

The representation of words using logical paradigms such as fuzzy logic, tensorial representations and other probabilistic approaches have been attempted before. In this section, we uncover some of these representations in detail.

Lee (1999) introduced measures of distributional similarity to improve the probability estimation for unseen occurrences. The measure of similarity of distributional word clusters was based on multiple measures including Euclidian distance, cosine distance, Jaccard's Coefficient, and asymmetric measures like α-skew divergence.

Bergmair (2011) used a fuzzy set theoretic view of features associated with word representations. While these features were not adopted from the vector space directly, it presents a unique perspective of entailment chains for reasoning tasks. Their analysis of inference using fuzzy representations provides interpretability in reasoning tasks.

Grefenstette (2013) presents a tenosrial calculus for word embeddings, which is based on compositional operators *which uses* vector representation of words to create a compositional distributional model of meaning. By providing a category-theoretic framework, the model creates an inherently compositional structure based on distributional word representations. However, they showed

that in this framework, quantifiers could not be expressed.

Herbelot and Vecchi (2015) refers to a notion of general formal semantics inferred from a distributional representation by creating relevant ontology based on the existing distribution. This mapping is therefore from a standard distributional model to a set-theoretic model, where dimensions are predicates and weights are generalised quantifiers.

Emerson and Copestake (2016, 2017) developed functional distributional semantics, which is a probabilistic framework based on model theory. The framework relies on differentiating and learning entities and predicates and their relations, on which Bayesian inference is performed. This representation is inherently compositional, context dependent representation.

3 Background: Fuzzy Sets and Fuzzy Logic

In this section, we provide a basic background of fuzzy sets including some fuzzy set operations, reinterpreting sets as tuples in a universe of finite elements and showing some set operations. We also cover the computation of fuzzy entropy as a Bernoulli random variable.

A fuzzy set is defined as a set with probabilistic set membership. Therefore, a fuzzy set is denoted as $A = \{(x, \mu_A(x)), x \in \Omega\}$, where x is an element of set A with a probability $\mu_A(x)$ such that $0 \le \mu_A \le 1$, and Ω is the universal set.

If our universe Ω is finite and of cardinality n, our notion of probabilistic set membership is constrained to a maximum n values. Therefore, each fuzzy set A can be represented as an n-tuple, with each member of the tuple $A[i]$ being the probability of the ith member of Ω. We can rewrite a fuzzy set as an n-tuple $A' = (\mu_{A'}(x), \forall x \in \Omega)$, such that $|A'| = |\Omega|$. In this representation, $A[i]$ is the probability of the ith member of the tuple A. We define some common set operations in terms of this representation as follows.

$$(A \cap B)[i] \equiv A[i] \times B[i] \quad \text{(set intersection)}$$
$$(A \cup B)[i] \equiv A[i] + B[i] - A[i] \times B[i] \text{ (set union)}$$
$$(A \sqcup B)[i] \equiv \max(1, \min(0, A[i] + B[i])) \text{ (disjoint union)}$$
$$(\neg A)[i] \equiv 1 - A[i] \quad \text{(complement)}$$
$$(A \setminus B)[i] \equiv A[i] - \min(A[i], B[i]) \quad \text{(set difference)}$$
$$(A \subseteq B) \equiv \forall x \in \Omega : \mu_A(x) \leq \mu_B(x) \text{ (set inclusion)}$$
$$|A| \equiv \sum_{i \in \Omega} \mu_A(i) \quad \text{(cardinality)}$$

The notion of entropy in fuzzy sets is an extrapolation of Shannon entropy from a single variable on the entire set. Formally, the fuzzy entropy of a set S is a measure of the uncertainty of the elements belonging to the set. The possibility of a member x belonging to the set S is a random variable X_i^S which is $true$ with probability (p_i^S) and $false$ with probability $(1 - p_i^S)$. Therefore, X_i^S is a Bernoulli random variable. In order to compute the entropy of a fuzzy set, we sum the entropy values of each X_i^S:

$$H(A) \equiv \sum_i H(X_i^A)$$
$$\equiv \sum_i -p_i^A \ln p_i^A - (1 - p_i^A) \ln(1 - p_i^A)$$
$$\equiv \sum_i -A[i] \ln A[i] - (1 - A[i]) \ln(1 - A[i])$$

This formulation will be useful in section 4.4 where we discuss two asymmetric measures of similarity, cross-entropy and K-L divergence, which can be seen as a natural extension of this formulation of fuzzy entropy.

4 Representation and Operations

In this section, we use the mathematical formulation above to reinterpret word embeddings. We first show how these word representations are created, then detail the interpretation of each of the set operations with some examples. We also look into some measures of similarity and their formulation in this framework. All examples in this section have been taken using the Google News Negative 300 vectors[1]. We used these gold standard vectors

---

[1] https://code.google.com/archive/p/word2vec/

4.1 Constructing the Tuple of Feature Probabilities

We start by converting the skip-gram negative sample word vectors into a tuple of feature probabilities. In order to construct a tuple of features representation in $\mathbb{R}^n$, we consider that the projection of a vector $\vec{v}$ onto a dimension i is a function of its probability of possessing the feature associated with that dimension. We compute the conversion from a word vector to a tuple of features by first exponentiating the projection of each vector along each direction, then averaging it over that feature for the entire vocabulary size, i.e. column-wise.

$$v_{exp}[i] \equiv \exp \vec{v}[i]$$
$$\hat{v}[i] \equiv \frac{v_{\exp}[i]}{\sum_{w \in \text{VOCAB}} \exp w_{\exp}[i]}$$

This normalization then produces a tuple of probabilities associated with each feature (corresponding to the dimensions of $\mathbb{R}^n$).

In line with our discussion from 3, this tuple of probabilities is akin to our representation of a fuzzy set. Let us consider the word v, and its corresponding n-dimensional word vector $\vec{v}$. The projection of $\vec{v}$ on a dimension i normalized (as shown above) to be interpreted as *if this dimension i were a property, what is probability that v would possess that property?*

In word2vec, words are distributed in a vector space of a particular dimensionality. Our representation attempts to provide some insight into how the arrangement of vectors provides insight into the properties they share. We do so by considering a function of the projection of a word vector onto a dimension and interpreting as a probability. This allows us an avenue to explore the relation between words in relation to the properties they share. It also allows us access to the entire arsenal of set operations, which are described below in section 4.2.

4.2 Operations on Feature Probabilities

Now that word vectors can be represented as tuples of feature probabilities, we can apply fuzzy set theoretic operations in order to ascertain the veracity of the implementation. We show qualitative examples of the set operations in this subsection, and the information they capture. Throughout this subsection, we follow the following notation: For any two words $w_1, w_2 \in$ VOCAB, $\hat{w}_1$ and $\hat{w}_2$ represents

$\hat{R}$	$\vec{R}$	$\hat{V}$	$\vec{V}$	$\hat{R} \cup \hat{V}$
risen	cashew	wavelengths	yellowish	flower
capita	risen	ultraviolet	whitish	red
peaked	soared	purple	aquamarine	stripes
declined	acuff	infrared	roans	flowers
increased	rafters	yellowish	bluish	green
rises	equalled	pigment	greenish	garlands

Table 1: An example of feature union. `Rose` is represented by R and `Violet` by V. We see here that while the word rose and violet have different meanings and senses, the union $R \cup V$ captures the sense of the flower as well as of colours, which are the senses common to these two words. We list words closest to the given word in the table. Closeness measured by cosine similarity for word2vec and cross-entropy-similarity for our vectors.

those words using our representation, while $\vec{w}_1$ and $\vec{w}_2$ are the word2vec vectors of those words.

Feature Union, Intersection and Difference In section 3, we showed the formulation of fuzzy set operations, assuming a finite universe of elements. As we saw in section 4.1, considering each dimension as a feature allows us to reinterpret word vectors as tuples of feature probabilities. Therefore, we can use the fuzzy set theoretic operations on this reinterpretation of fuzzy sets. For convenience, these operations have been called feature union, intersection and difference.

Intuitively, the feature intersection of words $\hat{w}_1$ and $\hat{w}_2$ should give us that word $\hat{w}_{1 \cap 2}$ which has the features common between the two words; an example of which is given in table 1. Similarly, the feature union $\hat{w}_{1 \cup 2} \simeq \hat{w}_1 \cup \hat{w}_2$ which has the properties of both the words, normalized for those properties which are common between the two, and feature difference $\hat{w}_{1 \setminus 2} \simeq \hat{w}_1 \setminus \hat{w}_2$ is that word which is similar to w_1 without the features of w_2. Examples of feature intersection and feature difference are shown in table 2 and 3 respectively.

While feature union does not seem to have a word2vec analogue, we consider that feature intersection is analogous to vector addition, and feature difference as analogous to vector difference.

Feature Inclusion Feature inclusion is based on the subset relation of fuzzy sets. We aim to capture feature inclusion by determining if there exist two words w_1 and w_2 such that *all* the feature probabilities of $\hat{w}_1$ are less than that of $\hat{w}_2$, then $\hat{w}_2 \subseteq \hat{w}_1$. We find that feature inclusion is closely linked to hyponymy, which we will show in 5.3.

$\hat{C}$	$\hat{P}$	$\hat{C} \cap \hat{P}$
hardware	vested	cpu
graphics	purchasing	hardware
multitasking	capita	powerpc
console	exercise	machine
firewire	parity	multitasking
mainframe	veto	microcode
$\vec{C}$	$\vec{P}$	$\vec{C} + \vec{P}$
bioses	centralize	expandability
scummvm	veto	writable
hardware	decembrist	cpcs
imovie	exercised	reconfigure
writable	redistribution	backplane
console	devolving	oem

Table 2: An example of feature intersection with the possible word2vec analogue (vector addition). The word `computer` is represented by C and `power` by P. Note that power is also a decent example of polysemy, and we see that in the context of computers, the connotations of hardware and the CPU are the most accessible. We list words closest to the given word in the table. Closeness measured by cosine similarity for word2vec and cross-entropy-similarity for our vectors.

4.3 Interpreting Entropy

For a word represented using a tuple of feature probabilities, the notion of entropy is strongly tied to the notion of certainty (Xuecheng, 1992), i.e. with what certainty does this word possess or not possess this set of features? Formally, the fuzzy entropy of a set S is a measure of the uncertainty of elements belonging to the set. The possibility a member x_i belonging to S is a random variable X_i^S, which is `true` with probability p_i^S, `false` with probability $(1 - p_i^S)$. Thus, X_i^S is a Bernoulli random variable. So, to measure the fuzzy entropy of a set, we add up the entropy values of each of the X_i^S (MacKay and Mac Kay, 2003).

Intuitively, words with the highest entropy are those which have features which are equally likely to belong to them and to their complement, i.e. $\forall i \in \Omega, A[i] \simeq 1 - A[i]$. So words with high fuzzy entropy can occur only in two scenarios: (1) The words occur with very low frequency so their random initialization remained, or (2) The words occur around so many different word groups that their corresponding fuzzy sets have some probability of possessing most of the features.

Therefore, our representation of words as tuples of features can be used to isolate function words better than the more commonly considered notion of simply using frequency, as it identifies the information theoretic distribution of features based on the context the function word occurs in. Table

$\hat{F}$	$\hat{B}$	$\hat{F} \setminus \hat{B}$
french	isles	communaut
english	colonial	aise
france	subcontinent	langue
german	cinema	monet
spanish	boer	dictionnaire
british	canadians	gascon
$\vec{F}$	$\vec{B}$	$\vec{F} - \vec{B}$
french	scottish	ranjit
english	american	privatised
france	thatcherism	tardis
german	netherlands	molloy
spanish	hillier	isaacs
british	cukcs	raj

Table 3: An example of feature difference, along with a possible word2vec analogue (vector difference). `French` is represented by F and `British` by B. We see here that set difference capture french words from the dataset, while there does not seem to be any such correlation in the vector difference. We list words closest to the given word in the table. Closeness measured by cosine similarity for word2vec and cross-entropy-similarity for our vectors.

4 provides the top 15 function words by entropy, and the correspodingly ranked words by frequency. We see that frequency is clearly not a good enough measure to identify function words.

4.4 Similarity Measures

One of the most important notions in presenting a distributional word representation is its ability to capture similarity (Van der Plas and Tiedemann, 2006). Since we use and modify vector based word representations, we aim to preserve the "distribution" of the vector embeddings, while providing a more robust interpretation of similarity measures. With respect to similarity, we make two strong claims:

1. Representing words as a tuple of feature probabilities lends us an inherent notion of similarity. Feature difference provides this notion, as it estimates the difference between two words along each feature probability.

2. Our representation allows for an easy adoption of known similarity measures such as K-L divergence and cross-entropy.

Note that feature difference (based on fuzzy set difference), K-L divergence and cross-entropy are all asymmetric measures of similarity. As Nematzadeh et al. (2017) points out, human similarity judgements are inherently asymmetric in nature.

We would like to point out that while most methods of introducing asymmetric similarity measures in word2vec account for both the focus and context vector Asr et al. (2018) and provide the asymmetry by querying on this combination of focus and context representations of each word. Our representation, on the other hand, uses only the focus representations (which are a part of the word representations used for downstream task as well as any other intrinsic evaluation), and still provides an innately asymmetric notion of similarity.

K-L Divergence From a fuzzy set perspective, we measure similarity as an overlap of features. For this purpose, we exploit the notion of fuzzy information theory by comparing how close the probability distributions of the similar words are using a standard measure, Kullback-Leibler (K-L) divergence. K-L divergence is an asymmetric measure of similarity.

The K-L divergence of a distribution P from another distribution Q is defined in terms of loss of compression. Given data d which follows distribution P, the extra bits need to store it under the false assumption that the data d follows distribution Q is the K-L divergence between the distributions P and Q. In the fuzzy case, we can compute the KL divergence as:

$$D(S \parallel T) \equiv D\left(X_i^S \,\middle\|\, X_i^T \right) = \sum_i p_i^S \log \left(p_i^S / p_i^T \right)$$

We see in table 5 some qualitative examples of how K-L divergence shows the relation between two words (or phrases when composed using feature intersection as in the case of `north korea`). We exemplify Nematzadeh et al. (2017)'s human annotator judgement of the distance between China and North Korea, where human annotators considered "North Korea" to be very similar to "China," while the reverse relationship was rated as significantly less strong ("China" is not very similar to "North Korea").

Cross Entropy We also calculate the cross entropy between two words, as it can be used to determine the entropy associated with the similarity between two words. Ideally, by determining the "spread" of the similarity of features between two words, we can determine the features that allow two words to be similar, allowing a more interpretable notion of feature-wise relation.

and	the	in	one	which	to	however	*two*	for	*eight*
this	of	of	in	the	*zero*	to	is	a	for
as	and	only	a	also	*nine*	it	as	but	*s*

Table 4: On the left: Top 15 words with highest entropy with frequency ≥ 100 (note that all of them are function words). On the right: Top 15 words with the highest frequency. The non-function words have been emphasized for comparison.

Example 1	$D(ganges \,\|\, delta)$			6.3105
	$D(delta \,\|\, ganges)$			6.3040
Example 2	$D(north \cap korea \,\|\, china)$			1.02923
	$D(china \,\|\, north \cap korea)$			10.60665

Table 5: Examples of KL-divergence as an asymmetric measure of similarity. Lower is closer. We see here that the evaluation of North Korea as a concept being closer to China than vice versa can be observed by the use of K-L Divergence on column-wise normalization.

$\hat{N}$	$\hat{M}$	$\hat{G}$	$\hat{N} \cap \hat{M}$	$\hat{N} \cap \hat{G}$
nobility	metal	bad	fusible	good
isotope	fusible	manners	unreactive	dharma
fujwara	ductility	happiness	metalloids	morals
feudal	with	evil	ductility	virtue
clan	alnico	excellent	heavy	righteous
$\vec{N}$	$\vec{M}$	$\vec{G}$	$\vec{N} + \vec{M}$	$\vec{N} + \vec{G}$
noblest	trivalent	bad	fusible	gracious
auctoritas	carbides	natured	metals	virtuous
abies	metallic	humoured	sulfides	believeth
eightfold	corrodes	selfless	finntroll	savages
vojt	alloying	gracious	rhodium	hedonist

Table 6: Polysemy of the word `noble`, in the context of the words `good` and `metal`. `noble` is represented by N, `metal` by M and `good` by G. We also provide the word2vec analogues of the same.

The cross-entropy of two distributions P and Q is a sum of the entropy of P and the K-L divergence between P and Q. In this sense, in captures both the *uncertainty in P*, as well as the distance from P to Q, to give us a general sense of the information theoretic difference between the concepts of P and Q. We use a generalized version of cross-entropy to fuzzy sets (Li, 2015), which is:

$$H(S, T) \equiv \sum_i H(X_i^S) + D(X_i^S \,\|\, X_i^T)$$

Feature representations which on comparison provide high cross entropy imply a more distributed feature space. Therefore, provided the right words to compute cross entropy, it could be possible to extract various features common (or associated) with a large group of words, lending some insight into how a single surface form (and its representation) can capture the distribution associated with different senses. Here, we use cross-entropy as a measure of polysemy, and isolate polysemous words based on context. We provide an example of capturing polysemy using composition by feature intersection in table 6.

We can see that the words which are most similar to `noble` are a combination of words from many senses, which provides some perspective into its distribution, . Indeed, it has an entropy value of 6.2765[2].

4.5 Constructing Analogy

Finally, we construct the notion of analogy in our representation of a word as a tuple of features. Word analogy is usually represented as a problem

where given a pairing $(a : b)$, and a prior x, we are asked to compute an unknown word $y_?$ such that $a : b :: x : y_?$. In the vector space model, analogy is computed based on vector distances. We find that this training mechanism does not have a consistent interpretation beyond evaluation. This is because normalization of vectors *performed only during inference, not during training*. Thus, computing analogy in terms of vector distances provides little insight into the distribution of vectors or to the notion of the length of the word vectors, which seems to be essential to analogy computation using vector operations

In using a fuzzy set theoretic representation, vector projections are inherently normalized, making them feature dense. This allows us to compute analogies much better in lower dimension spaces. We consider analogy to be an operation involving union and set difference. Word analogy is computed as follows:

$$a : b :: x : y_?$$
$$y_? = b - a + x \implies y_? = (b + x) - a$$
$$y = (b \sqcup x) \setminus a \quad \text{(Set-theoretic interpretation)}$$

Notice that this form of word analogy can be "derived" from the vector formula by re-arrangement. We use non-disjoint set union so that the common features are not eliminated, but the values

Word 1	Word 2	Word 3	word2vec	Our representation
bacteria	tuberculosis	virus	polio	hiv
cold	freezing	hot	evaporates	boiling
ds	nintendo	dreamcast	playstation	sega
pool	billiards	karate	taekwondo	judo

Table 7: Examples of analogy compared to the analogy in word2vec. We see here that the comparisons constructed by feature representations are similar to those given by the standard word vectors.

are clipped at $(0, 1]$ so that the fuzzy representation is consistent. Analogical reasoning is based on the common features between the word representations, and conflates multiple types of relations such as synonymy, hypernymy and causal relations (Chen et al., 2017). Using fuzzy set theoretic representations, we can also provide a context for the analogy, effectively reconstructing analogous reasoning to account for the type of relation from a lexical semantic perspective.

Some examples of word analogy based are presented in table 7.

5 Experiments and Results

In this section, we present our experiments and their results in various domains including similarity, analogy, function word detection, polysemy detection, lexical entailment and compositionality. All the experiments have been conducted on established datasets.

5.1 Similarity and Analogy

Similarity and analogy are the most popular intrinsic evaluation mechanisms for word representations (Mikolov et al., 2013a). Therefore, to evaluate our representations, the first tasks we show are similarity and analogy. For similarity computations, we use the SimLex corpus (Hill et al., 2015) for training and testing at different dimensions For word analogy, we use the MSR Word Relatedness Test (Mikolov et al., 2013c). We compare it to the vector representation of words for different dimensions.

5.1.1 Similarity

Our scores our compared to the word2vec scores of similarity using the Spearman rank correlation coefficient (Spearman, 1987), which is a ratio of the covariances and standard deviations of the inputs being compared.

As shown in table 8, using our representation, similarity is *slightly* better represented according to the SimLex corpus. We show similarity on both the asymmetric measures of similarity for our repre-

Dims.	word2vec	Our Representation	
		K-L Divergence	Cross-Entropy
20	0.2478	0.2690	**0.2744**
50	0.2916	0.2966	**0.2981**
100	0.2960	0.3124	**0.3206**
200	0.3259	0.3253	**0.3298**

Table 8: Similarity scores on the SimLex-999 dataset (Hill et al., 2015), for various dimension sizes (Dims.). The scores are provided according to the Spearman Correlation to incorporate higher precision.

Category		word2vec		Our representation	
		50	100	50	100
Capital Common Countries		21.94	37.55	**39.13**	**47.23**
Capital World		13.02	20.10	**27.30**	**26.54**
Currency		12.24	18.60	**25.27**	**24.90**
City-State		10.38	16.70	**23.24**	**23.51**
Family		10.61	17.34	**23.67**	**23.88**
Adjective-Adverb	Syntactic	4.74	3.23	**7.26**	**3.83**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	9.92	15.68	**21.73**	**21.52**
Opposite	Syntactic	4.06	3.66	**7.61**	**4.92**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	9.36	14.73	**20.60**	**20.26**
Comparative	Syntactic	8.86	12.63	**16.88**	**15.39**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	10.10	15.96	**21.67**	**21.39**
Superlative	Syntactic	7.59	11.30	**14.32**	**13.36**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	9.54	15.20	**20.35**	**20.15**
Present-Participle	Syntactic	7.51	10.96	**14.31**	**13.14**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	9.34	14.73	**19.84**	**19.49**
Nationality	Syntactic	12.51	19.07	**21.64**	**21.96**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	11.51	18.16	**22.71**	**22.97**
Past Tense	Syntactic	11.65	17.09	**20.43**	**19.76**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	11.16	17.21	**21.96**	**27.72**
Plural	Syntactic	11.76	17.23	**20.53**	**19.89**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	11.26	17.28	**21.90**	**21.64**
Plural Verbs	Syntactic	11.36	16.60	**19.88**	**19.46**
	Semantic	10.61	17.34	**23.67**	**23.88**
	Overall	11.05	16.91	**21.46**	**21.30**

Table 9: Comparison of Analogies between word2vec and our representation for 50 and 100 dimensions (Dims.). For the first five, only overall accuracy is shown as overall accuracy is the same as semantic accuracy (as there is no syntactic accuracy measure). For all the others, we present, syntactic, semantic and overall accuracy as well. We see here that we outperform word2vec on every single metric.

sentation, K-L divergence as well as cross-entropy. We see that cross-entropy performs better than K-L Divergence. While the similarity scores are generally higher, we see a reduction in the degree of similarity beyond 100 dimension vectors (features).

5.1.2 Analogy

For analogy, we see that our model outperforms word2vec at both 50 and 100 dimensions. We see that at lower dimension sizes, our normalized feature representation captures significantly more syntactic and semantic information than its vector counterpart. We conjecture that this can primarily be attributed to the fact that constructing feature probabilities provides more information about the

top n words	word2vec	Our Representation
15	10	**15**
30	21	**30**
50	39	**47**

Table 10: Function word detection using entropy (in our representation) and by frequency in word2vec. We see that we consistently detect more function words than word2vec, based on the 176 function word list released by Nation (2016). The metric is *number of words*, i.e. the number of words chosen by frequency for word2vec and entropy for our representation

Dims.	Metric	word2vec	Our Representation
50	Spearman	0.3946	**0.4117**
	Pearson	0.4058	**0.4081**
100	Spearman	0.4646	**0.4912**
	Pearson	0.4457	**0.4803**
200	Spearman	0.4479	**0.4549**
	Pearson	**0.4163**	0.4091

Table 11: Results for compositionality of word embeddings for nominal compounds for various dimensions (Dims.). We see that almost across the board, we perform better, however, for the Pearson correlation metric, at 200 dimensions, we find that word2vec has a better representation of rank by frequency for nominal compounds.

common (and distinct) "concepts" which are shared between two words.

Since feature representations are inherently fuzzy sets, lower dimension sizes provide a more reliable probability distribution, which becomes more and more sparse as the dimensionality of the vectors increases (i.e. number of features rise). Therefore, we notice that the increase in feature probabilities is a lot more for 50 dimensions than it is for 100.

5.2 Function Word Detection

As mentioned in section 4.3, we use entropy as a measure of detecting function words for the standard GoogleNews-300 negative sampling dataset[3]. In order to quantitatively evaluate the detection of function words, we choose the top n words in our representation ordered by entropy with a frequency ≥ 100, and compare it to the top n words ordered by frequency from word2vec; n being 15, 30 and 50. We compare the number of function words in both in table 10. The list of function words is derived from Nation (2016).

5.3 Compositionality

Finally, we evaluate the compositionality of word embeddings. Mikolov et al. (2013b) claims that word embeddings in vector spaces possess additive compositionality, i.e. by vector addition, semantic phrases such as compounds can be well represented. We claim that our representation in fact captures the semantics of phrases by performing a literal combination of the features of the head and modifier word, therefore providing a more robust representation of phrases.

We use the English nominal compound phrases from Ramisch et al. (2016). An initial set of experiments on nominal compounds using word2vec have been done before (Cordeiro et al., 2016), where it

was shown to be a fairly difficult task for modern non-contextual word embeddings. In order to analyse nominal compounds, we adjust our similarity metric to account for asymmetry in the similarity between the head-word and the modifier, and vice versa. We report performance on two metrics, the Spearman correlation (Spearman, 1987) and Pearson correlation (Pearson, 1920).

The results are shown in table 11. The difference in scores for the Pearson and Spearman rank correlation show that word2vec at higher dimensions better represents the rank of words (by frequency), but at lower dimensions, the feature probability representation has a better analysis of both rank by frequency, and its correlation with similarity of words with a nominal compound. Despite this, we show a higher Spearman correlation coefficient at 200 dimesions as well, as we capture non-linear relations.

5.4 Dimensionality Analysis and Feature Representations

In this subsection, we provide some interpretation of the results above, and examine the effect of scaling dimensions to the feature representation. As seen here, the evaluation has been done on smaller dimension sizes of 50 and 100, and we see that our representation can be used for a slightly larger range of tasks from the perspective of intrinsic evaluations. However, the results of quantitative analogy for higher dimensions have been observed to be lower for fuzzy representations rather than the word2vec negative-sampling word vectors.

We see that the representation we propose does not scale well as dimensions increase. This is because our representation relies on the distribution of probability mass per feature (dimension) across all the words. Therefore, increasing the dimension-

---

[3]https://code.google.com/archive/p/word2vec/

ality of the word vectors used makes the representation that much more sparse.

6 Conclusion

In this paper, we presented a reinterpretation of distributional semantics. We performed a column-wise normalization on word vectors, such that each value in this normalized representation represented the probability of the word possessing a feature that corresponded to each dimension. This provides us a representation of each word as a tuple of feature probabilities. We find that this representation can be seen as a fuzzy set, with each probability being the function of the projection of the original word vector on a dimension.

Considering word vectors as fuzzy sets allows us access to set operations such as union, intersection and difference. In our modification, these operations provide the product, disjoint sum and difference of the word representations, feature wise. Using qualitative examples, we show that our representation naturally captures an asymmetric notion of similarity using feature difference, from which known asymmetric measures can be easily constructed, such as Cross Entropy and K-L Divergence.

We qualitatively show how our model accounts for polysemy, while showing quantitative proofs of our representation's performance at lower dimensions in similarity, analogy, compositionality and function word detection. We hypothesize that lower dimensions are more suited for our representation as sparsity increases with higher dimensions, so the significance of feature probabilities reduces. This sparsity causes a diffusion of the probabilities across multiple features.

Through this work, we aim to provide some insights into interpreting word representations by showing one possible perspective and explanation of the lengths and projections of word embeddings in the vector space. These feature representations can be adapted for basic neural models, allowing the use of feature based representations at lower dimensions for downstream tasks.

Acknowledgements

References

Fatemeh Torabi Asr, Robert Zinkov, and Michael Jones. 2018. Querying word embeddings for similarity and relatedness. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 675–684.

Richard Bergmair. 2011. *Monte Carlo Semantics: Robust inference and logical pattern processing with Natural Language text*. Ph.D. thesis, University of Cambridge.

Jose Camacho-Collados and Mohammad Taher Pilehvar. 2018. From word to sense embeddings: A survey on vector representations of meaning. *Journal of Artificial Intelligence Research*, 63:743–788.

Dawn Chen, Joshua C Peterson, and Thomas L Griffiths. 2017. Evaluating vector-space models of analogy. *arXiv preprint arXiv:1705.04416*.

Silvio Cordeiro, Carlos Ramisch, Marco Idiart, and Aline Villavicencio. 2016. Predicting the compositionality of nominal compounds: Giving word embeddings a hard time. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1986–1997, Berlin, Germany. Association for Computational Linguistics.

Guy Emerson and Ann Copestake. 2016. Functional distributional semantics. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 40–52.

Guy Emerson and Ann Copestake. 2017. Semantic composition via probabilistic model theory. In *IWCS 2017-12th International Conference on Computational Semantics-Long papers*.

John R Firth. 1957. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*.

Jean Mark Gawron. 2014. Improving sparse word similarity models with asymmetric measures. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 296–301.

Yoav Goldberg and Omer Levy. 2014. word2vec explained: deriving mikolov et al.'s negative-sampling word-embedding method. *arXiv preprint arXiv:1402.3722*.

Edward Grefenstette. 2013. Towards a formal distributional semantics: Simulating logical calculi with tensors. *arXiv preprint arXiv:1304.5823*.

James Henderson and Diana Nicoleta Popa. 2016. A vector space for distributional semantics for entailment. *arXiv preprint arXiv:1607.03780*.

Aurélie Herbelot and Eva Maria Vecchi. 2015. Building a shared world: Mapping distributional to model-theoretic semantic spaces. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 22–32.

Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.

Lillian Lee. 1999. Measures of distributional similarity. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 25–32.

Jiwei Li and Dan Jurafsky. 2015. Do multi-sense embeddings improve natural language understanding? *arXiv preprint arXiv:1506.01070*.

Xiang Li. 2015. Fuzzy cross-entropy. *Journal of Uncertainty Analysis and Applications*, 3(1):2.

David JC MacKay and David JC Mac Kay. 2003. *Information theory, inference and learning algorithms*. Cambridge university press.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.

Tomáš Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013c. Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: Human language technologies*, pages 746–751.

Ian Stephen Paul Nation. 2016. *Making and using word lists for language learning and testing*. John Benjamins Publishing Company.

Aida Nematzadeh, Stephan C Meylan, and Thomas L Griffiths. 2017. Evaluating vector-space models of word representation, or, the unreasonable effectiveness of counting words near other words. In *CogSci*.

Karl Pearson. 1920. Notes on the history of correlation. *Biometrika*, 13(1):25–45.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.

Lonneke Van der Plas and Jörg Tiedemann. 2006. Finding synonyms using automatic word alignment and measures of distributional similarity. In *Proceedings of the COLING/ACL on Main conference poster sessions*, pages 866–873. Association for Computational Linguistics.

Carlos Ramisch, Silvio Cordeiro, Leonardo Zilio, Marco Idiart, and Aline Villavicencio. 2016. How naked is the naked truth? a multilingual lexicon of nominal compound compositionality. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 156–161, Berlin, Germany. Association for Computational Linguistics.

Charles Spearman. 1987. The proof and measurement of association between two things. *The American journal of psychology*, 100(3/4):441–471.

Amos Tversky. 1977. Features of similarity. *Psychological review*, 84(4):327.

Luke Vilnis and Andrew McCallum. 2014. Word representations via gaussian embedding. *arXiv preprint arXiv:1412.6623*.

Liu Xuecheng. 1992. Entropy, distance measure and similarity measure of fuzzy sets and their relations. *Fuzzy sets and systems*, 52(3):305–318.