

Distributional Inclusion Vector Embedding for Unsupervised Hypernymy Detection

Haw-Shiuan Chang¹, ZiYun Wang², Luke Vilnis¹, Andrew McCallum¹

¹University of Massachusetts, Amherst, USA

²Tsinghua University, Beijing, China

hschang@cs.umass.edu, wang-zy14@mails.tsinghua.edu.cn,
{luke, mccallum}@cs.umass.edu

Abstract

Modeling hypernymy, such as poodle is-a dog, is an important generalization aid to many NLP tasks, such as entailment, coreference, relation extraction, and question answering. Supervised learning from labeled hypernym sources, such as WordNet, limits the coverage of these models, which can be addressed by learning hypernyms from unlabeled text. Existing unsupervised methods either do not scale to large vocabularies or yield unacceptably poor accuracy. This paper introduces *distributional inclusion vector embedding (DIVE)*, a simple-to-implement unsupervised method of hypernym discovery via per-word non-negative vector embeddings which preserve the inclusion property of word contexts in a low-dimensional and interpretable space. In experimental evaluations more comprehensive than any previous literature of which we are aware—evaluating on 11 datasets using multiple existing as well as newly proposed scoring functions—we find that our method provides up to double the precision of previous unsupervised embeddings, and the highest average performance, using a much more compact word representation, and yielding many new state-of-the-art results.

1 Introduction

Numerous applications benefit from compactly representing context distributions, which assign meaning to objects under the rubric of *distributional semantics*. In natural language processing, distributional semantics has long been used to assign meanings to words (that is, to *lexemes* in the dictionary, not individual instances of word tokens). The meaning of a word in the distributional sense is often taken to be the set of textual contexts (nearby tokens) in which that word appears, represented as a large sparse bag of words (SBOW). Without any supervision,

Word2Vec (Mikolov et al., 2013), among other approaches based on matrix factorization (Levy et al., 2015a), successfully compress the SBOW into a much lower dimensional embedding space, increasing the scalability and applicability of the embeddings while preserving (or even improving) the correlation of geometric embedding similarities with human word similarity judgments.

While embedding models have achieved impressive results, context distributions capture more semantic information than just word similarity. The *distributional inclusion hypothesis (DIH)* (Weeds and Weir, 2003; Geffet and Dagan, 2005; Cimiano et al., 2005) posits that the context set of a word tends to be a subset of the contexts of its hypernyms. For a concrete example, most adjectives that can be applied to poodle can also be applied to dog, because dog is a hypernym of poodle (e.g. both can be obedient). However, the converse is not necessarily true — a dog can be straight-haired but a poodle cannot. Therefore, dog tends to have a broader context set than poodle. Many asymmetric scoring functions comparing SBOW features based on DIH have been developed for hypernymy detection (Weeds and Weir, 2003; Geffet and Dagan, 2005; Shwartz et al., 2017).

Hypernymy detection plays a key role in many challenging NLP tasks, such as textual entailment (Sammons et al., 2011), coreference (Ponzetto and Strube, 2006), relation extraction (Demeester et al., 2016) and question answering (Huang et al., 2008). Leveraging the variety of contexts and inclusion properties in context distributions can greatly increase the ability to discover taxonomic structure among words (Shwartz et al., 2017). The inability to preserve these features limits the semantic representation power and downstream applicability of some popular unsupervised learning approaches such as Word2Vec.

Several recently proposed methods aim to en-

code hypernym relations between words in dense embeddings, such as Gaussian embedding (Vilnis and McCallum, 2015; Athiwaratkun and Wilson, 2017), Boolean Distributional Semantic Model (Kruszewski et al., 2015), order embedding (Vendrov et al., 2016), H-feature detector (Roller and Erk, 2016), HyperVec (Nguyen et al., 2017), dual tensor (Glavaš and Ponzetto, 2017), Poincaré embedding (Nickel and Kiela, 2017), and LEAR (Vulić and Mrkšić, 2017). However, the methods focus on supervised or semi-supervised settings where a massive amount of hypernym annotations are available (Vendrov et al., 2016; Roller and Erk, 2016; Nguyen et al., 2017; Glavaš and Ponzetto, 2017; Vulić and Mrkšić, 2017), do not learn from raw text (Nickel and Kiela, 2017) or lack comprehensive experiments on the hypernym detection task (Vilnis and McCallum, 2015; Athiwaratkun and Wilson, 2017).

Recent studies (Levy et al., 2015b; Shwartz et al., 2017) have underscored the difficulty of generalizing supervised hypernymy annotations to unseen pairs — classifiers often effectively memorize prototypical hypernyms (‘general’ words) and ignore relations between words. These findings motivate us to develop more accurate and scalable unsupervised embeddings to detect hypernymy and propose several scoring functions to analyze the embeddings from different perspectives.

1.1 Contributions

- A novel unsupervised low-dimensional embedding method via performing non-negative matrix factorization (NMF) on a weighted PMI matrix, which can be efficiently optimized using modified skip-grams.
- Theoretical and qualitative analysis illustrate that the proposed embedding can intuitively and interpretably preserve inclusion relations among word contexts.
- Extensive experiments on 11 hypernym detection datasets demonstrate that the learned embeddings dominate previous low-dimensional unsupervised embedding approaches, achieving similar or better performance than SBOW, on both existing and newly proposed asymmetric scoring functions, while requiring much less memory and compute.

2 Method

The *distributional inclusion hypothesis* (DIH) suggests that the context set of a hypernym tends to contain the context set of its hyponyms. When representing a word as the counts of contextual co-occurrences, the count in every dimension of hypernym y tends to be larger than or equal to the corresponding count of its hyponym x :

$$x \preceq y \iff \forall c \in V, \#(x, c) \leq \#(y, c), \quad (1)$$

where $x \preceq y$ means y is a hypernym of x , V is the set of vocabulary, and $\#(x, c)$ indicates the number of times that word x and its context word c co-occur in a small window with size $|W|$ in the corpus of interest D . Notice that the concept of DIH could be applied to different context word representations. For example, Geffet and Dagan (2005) represent each word by the set of its co-occurred context words while discarding their counts. In this study, we define the inclusion property based on counts of context words in (1) because the counts are an effective and noise-robust feature for the hypernymy detection using only the context distribution of words (Clarke, 2009; Vulić et al., 2016; Shwartz et al., 2017).

Our goal is to produce lower-dimensional embeddings preserving the inclusion property that the embedding of hypernym y is larger than or equal to the embedding of its hyponym x in every dimension. Formally, the desired property can be written as

$$x \preceq y \iff \mathbf{x}[i] \leq \mathbf{y}[i], \forall i \in \{1, \dots, L\}, \quad (2)$$

where L is number of dimensions in the embedding space. We add additional non-negativity constraints, i.e. $\mathbf{x}[i] \geq 0, \mathbf{y}[i] \geq 0, \forall i$, in order to increase the interpretability of the embeddings (the reason will be explained later in this section).

This is a challenging task. In reality, there are a lot of noise and systematic biases that cause the violation of DIH in Equation (1) (i.e. $\#(x, c) > \#(y, c)$ for some neighboring word c), but the general trend can be discovered by processing thousands of neighboring words in SBOW together (Shwartz et al., 2017). After the compression, the same trend has to be estimated in a much smaller embedding space which discards most of the information in SBOW, so it is not surprising to see most of the unsupervised hypernymy detection studies focus on SBOW (Shwartz et al., 2017)

and the existing unsupervised embedding methods like Gaussian embedding have degraded accuracy (Vulić et al., 2016).

2.1 Inclusion Preserving Matrix Factorization

Popular methods of unsupervised word embedding are usually based on matrix factorization (Levy et al., 2015a). The approaches first compute a co-occurrence statistic between the w th word and the c th context word as the (w, c) th element of the matrix $M[w, c]$. Next, the matrix M is factorized such that $M[w, c] \approx \mathbf{w}^T \mathbf{c}$, where \mathbf{w} is the low dimension embedding of w th word and \mathbf{c} is the c th context embedding.

The statistic in $M[w, c]$ is usually related to pointwise mutual information (Levy et al., 2015a): $PMI(w, c) = \log(\frac{P(w, c)}{P(w) \cdot P(c)})$, where $P(w, c) = \frac{\#(w, c)}{|D|}$, $|D| = \sum_{w \in V} \sum_{c \in V} \#(w, c)$ is number of co-occurrence word pairs in the corpus, $P(w) = \frac{\#(w)}{|D|}$, $\#(w) = \sum_{c \in V} \#(w, c)$ is the frequency of the word w times the window size $|W|$, and similarly for $P(c)$. For example, $M[w, c]$ could be set as positive PMI (PPMI), $\max(PMI(w, c), 0)$, or shifted PMI, $PMI(w, c) - \log(k')$, which (Levy and Goldberg, 2014) demonstrate is connected to skip-grams with negative sampling (SGNS).

Intuitively, since $M[w, c] \approx \mathbf{w}^T \mathbf{c}$, larger embedding values of \mathbf{w} at every dimension seems to imply larger $\mathbf{w}^T \mathbf{c}$, larger $M[w, c]$, larger $PMI(w, c)$, and thus larger co-occurrence count $\#(w, c)$. However, the derivation has two flaws: (1) \mathbf{c} could contain negative values and (2) lower $\#(w, c)$ could still lead to larger $PMI(w, c)$ as long as the $\#(w)$ is small enough.

To preserve DIH, we propose a novel word embedding method, *distributional inclusion vector embedding (DIVE)*, which fixes the two flaws by performing non-negative factorization (NMF) (Lee and Seung, 2001) on the matrix M , where $M[w, c] =$

$$\log\left(\frac{P(w, c)}{P(w) \cdot P(c)} \cdot \frac{\#(w)}{k_I \cdot Z}\right) = \log\left(\frac{\#(w, c)|V|}{\#(c)k_I}\right), \quad (3)$$

where k_I is a constant which shifts PMI value like SGNS, $Z = \frac{|D|}{|V|}$ is the average word frequency, and $|V|$ is the vocabulary size. We call this weighting term $\frac{\#(w)}{Z}$ *inclusion shift*.

After applying the non-negativity constraint and inclusion shift, the inclusion property in DIVE

(i.e. Equation (2)) implies that Equation (1) (DIH) holds if the matrix is reconstructed perfectly. The derivation is simple: If the embedding of hypernym \mathbf{y} is greater than or equal to the embedding of its hyponym \mathbf{x} in every dimension ($\mathbf{x}[i] \leq \mathbf{y}[i], \forall i$), $\mathbf{x}^T \mathbf{c} \leq \mathbf{y}^T \mathbf{c}$ since context vector \mathbf{c} is non-negative. Then, $M[x, c] \leq M[y, c]$ tends to be true because $\mathbf{w}^T \mathbf{c} \approx M[w, c]$. This leads to $\#(x, c) \leq \#(y, c)$ because $M[w, c] = \log(\frac{\#(w, c)|V|}{\#(c)k_I})$ and only $\#(w, c)$ changes with w .

2.2 Optimization

Due to its appealing scalability properties during training time (Levy et al., 2015a), we optimize our embedding based on the skip-gram with negative sampling (SGNS) (Mikolov et al., 2013). The objective function of SGNS is

$$l_{SGNS} = \sum_{w \in V} \sum_{c \in V} \#(w, c) \log \sigma(\mathbf{w}^T \mathbf{c}) + \sum_{w \in V} k' \sum_{c \in V} \#(w, c) \mathbb{E}_{c_N \sim P_D} [\log \sigma(-\mathbf{w}^T \mathbf{c}_N)], \quad (4)$$

where $\mathbf{w} \in \mathbb{R}$, $\mathbf{c} \in \mathbb{R}$, $\mathbf{c}_N \in \mathbb{R}$, σ is the logistic sigmoid function, and k' is a constant hyper-parameter indicating the ratio between positive and negative samples.

Levy and Goldberg (2014) demonstrate SGNS is equivalent to factorizing a shifted PMI matrix M' , where $M'[w, c] = \log(\frac{P(w, c)}{P(w) \cdot P(c)} \cdot \frac{1}{k'})$. By setting $k' = \frac{k_I \cdot Z}{\#(w)}$ and applying non-negativity constraints to the embeddings, DIVE can be optimized using the similar objective function:

$$l_{DIVE} = \sum_{w \in V} \sum_{c \in V} \#(w, c) \log \sigma(\mathbf{w}^T \mathbf{c}) + k_I \sum_{w \in V} \frac{Z}{\#(w)} \sum_{c \in V} \#(w, c) \mathbb{E}_{c_N \sim P_D} [\log \sigma(-\mathbf{w}^T \mathbf{c}_N)], \quad (5)$$

where $\mathbf{w} \geq 0$, $\mathbf{c} \geq 0$, $\mathbf{c}_N \geq 0$, and k_I is a constant hyper-parameter. P_D is the distribution of negative samples, which we set to be the corpus word frequency distribution (not reducing the probability of drawing frequent words like SGNS) in this paper. Equation (5) is optimized by ADAM (Kingma and Ba, 2015), a variant of stochastic gradient descent (SGD). The non-negativity constraint is implemented by projection (Polyak, 1969) (i.e. clipping any embedding which crosses the zero boundary after an update).

The optimization process provides an alternative angle to explain how DIVE preserves DIH.




Figure 1: The embedding of the words rodent and mammal trained by the co-occurrence statistics of context words using DIVE. The index of dimensions is sorted by the embedding values of mammal and values smaller than 0.1 are neglected. The top 5 words (sorted by its embedding value of the dimension) tend to be more general or more representative on the topic than the top 51-105 words.

The gradients for the word embedding \mathbf{w} is

$$\frac{dl_{DIVE}}{d\mathbf{w}} = \sum_{c \in V} \#(w, c)(1 - \sigma(\mathbf{w}^T \mathbf{c}))\mathbf{c} - k_I \sum_{c_N \in V} \frac{\#(c_N)}{|V|} \sigma(\mathbf{w}^T \mathbf{c}_N)\mathbf{c}_N. \quad (6)$$

Assume hyponym x and hypernym y satisfy DIH in Equation (1) and the embeddings \mathbf{x} and \mathbf{y} are the same at some point during the gradient ascent. At this point, the gradients coming from negative sampling (the second term) decrease the same amount of embedding values for both x and y . However, the embedding of hypernym y would get higher or equal positive gradients from the first term than x in every dimension because $\#(x, c) \leq \#(y, c)$. This means Equation (1) tends to imply Equation (2) because the hypernym has larger gradients everywhere in the embedding space.

Combining the analysis from the matrix factorization viewpoint, DIH in Equation (1) is approximately equivalent to the inclusion property in DIVE (i.e. Equation (2)).

2.3 PMI Filtering

For a frequent target word, there must be many neighboring words that incidentally appear near the target word without being semantically meaningful, especially when a large context window size is used. The unrelated context words cause noise in both the word vector and the context vector of DIVE. We address this issue by filtering out context words c for each target word w when the PMI of the co-occurring words is too small (i.e. $\log(\frac{P(w, c)}{P(w) \cdot P(c)}) < \log(k_f)$). That is, we set

$\#(w, c) = 0$ in the objective function. This pre-processing step is similar to computing PPMI in SBOW (Bullinaria and Levy, 2007), where low PMI co-occurrences are removed from SBOW.

2.4 Interpretability

After applying the non-negativity constraint, we observe that each latent factor in the embedding is interpretable as previous findings suggest (Pauca et al., 2004; Murphy et al., 2012) (i.e. each dimension roughly corresponds to a topic). Furthermore, DIH suggests that a general word appears in more diverse contexts/topics. By preserving DIH using inclusion shift, the embedding of a general word (i.e. hypernym of many other words) tends to have larger values in these dimensions (topics). This gives rise to a natural and intuitive interpretation of our word embeddings: the word embeddings can be seen as unnormalized probability distributions over topics. In Figure 1, we visualize the unnormalized topical distribution of two words, rodent and mammal, as an example. Since rodent is a kind of mammal, the embedding (i.e. unnormalized topical distribution) of mammal includes the embedding of rodent when DIH holds. More examples are illustrated in our supplementary materials.

3 Unsupervised Embedding Comparison

In this section, we compare DIVE with other unsupervised hypernym detection methods. In this paper, unsupervised approaches refer to the methods that only train on plaintext corpus without using any hypernymy or lexicon annotation.

Dataset	BLESS	EVALution	LenciBenotto	Weeds	Medical	LEDS
Random	5.3	26.6	41.2	51.4	8.5	50.5
Word2Vec + C	9.2	25.4	40.8	51.6	11.2	71.8
GE + C	10.5	26.7	43.3	52.0	14.9	69.7
GE + KL	7.6	29.6	45.1	51.3	15.7	64.6 (80 ³)
DIVE + C- Δ S	16.3	33.0	50.4	65.5	19.2	83.5

Dataset	TM14	Kotlerman 2010	HypeNet	WordNet	Avg (10 datasets)	HyperLex
Random	52.0	30.8	24.5	55.2	23.2	0
Word2Vec + C	52.1	39.5	20.7	63.0	25.3	16.3
GE + C	53.9	36.0	21.6	58.2	26.1	16.4
GE + KL	52.0	39.4	23.7	54.4	25.9	9.6 (20.6 ³)
DIVE + C- Δ S	57.2	36.6	32.0	60.9	32.7	32.8

Table 1: Comparison with other unsupervised embedding methods. The scores are AP@all (%) for the first 10 datasets and Spearman ρ (%) for HyperLex. Avg (10 datasets) shows the micro-average AP of all datasets except HyperLex. Word2Vec+C scores word pairs using cosine similarity on skip-grams. GE+C and GE+KL compute cosine similarity and negative KL divergence on Gaussian embedding, respectively.

3.1 Experiment Setup

The embeddings are tested on 11 datasets. The first 4 datasets come from the recent review of Shwartz et al. (2017)¹: BLESS (Baroni and Lenci, 2011), EVALution (Santus et al., 2015), Lenci/Benotto (Benotto, 2015), and Weeds (Weeds et al., 2014). The next 4 datasets are downloaded from the code repository of the H-feature detector (Roller and Erk, 2016)²: Medical (i.e., Levy 2014) (Levy et al., 2014), LE DS (also referred to as ENTAILMENT or Baroni 2012) (Baroni et al., 2012), TM14 (i.e., Turney 2014) (Turney and Mohammad, 2015), and Kotlerman 2010 (Kotlerman et al., 2010). In addition, the performance on the test set of HypeNet (Shwartz et al., 2016) (using the random train/test split), the test set of WordNet (Vendrov et al., 2016), and all pairs in HyperLex (Vulić et al., 2016) are also evaluated.

The F1 and accuracy measurements are sometimes very similar even though the quality of prediction varies, so we adopted average precision, AP@all (Zhu, 2004) (equivalent to the area under the precision-recall curve when the constant interpolation is used), as the main evaluation metric. The HyperLex dataset has a continuous score on each candidate word pair, so we adopt Spearman rank coefficient ρ (Fieller et al., 1957) as suggested by the review study of Vulić et al. (2016). Any OOV (out-of-vocabulary) word encountered in the testing data is pushed to the bottom of the prediction list (effectively assuming the word pair does not have hypernym relation).

¹<https://github.com/vered1986/UnsupervisedHypernymy>

²<https://github.com/stephenroller/emnlp2016/>

We trained all methods on the first 51.2 million tokens of WaCkypedia corpus (Baroni et al., 2009) because DIH holds more often in this subset (i.e. SBOW works better) compared with that in the whole WaCkypedia corpus. The window size $|W|$ of DIVE and Gaussian embedding are set as 20 (left 10 words and right 10 words). The number of embedding dimensions in DIVE L is set to be 100. The other hyper-parameters of DIVE and Gaussian embedding are determined by the training set of HypeNet. Other experimental details are described in our supplementary materials.

3.2 Results

If a pair of words has hypernym relation, the words tend to be similar (sharing some context words) and the hypernym should be more general than the hyponym. Section 2.4 has shown that the embedding could be viewed as an unnormalized topic distribution of its context, so the embedding of hypernym should be similar to the embedding of its hyponym but having larger magnitude. As in HyperVec (Nguyen et al., 2017), we score the hypernym candidates by multiplying two factors corresponding to these properties. The $C \cdot \Delta S$ (i.e. the cosine similarity multiply the difference of summation) scoring function is defined as

$$C \cdot \Delta S(\mathbf{w}_q \rightarrow \mathbf{w}_p) = \frac{\mathbf{w}_q^T \mathbf{w}_p}{\|\mathbf{w}_q\|_2 \cdot \|\mathbf{w}_p\|_2} \cdot (\|\mathbf{w}_p\|_1 - \|\mathbf{w}_q\|_1), \quad (7)$$

where \mathbf{w}_p is the embedding of hypernym and \mathbf{w}_q is the embedding of hyponym.

As far as we know, Gaussian embedding (GE) (Vilnis and McCallum, 2015) is the state-of-the-art unsupervised embedding method which can capture the asymmetric relations between a hypernym and its hyponyms. Gaussian embedding

encodes the context distribution of each word as a multivariate Gaussian distribution, where the embeddings of hypernyms tend to have higher variance and overlap with the embedding of their hyponyms. In Table 1, we compare DIVE with Gaussian embedding³ using the code implemented by Athiwaratkun and Wilson (2017)⁴ and with word cosine similarity using skip-grams. The performances of random scores are also presented for reference. As we can see, DIVE is usually significantly better than other unsupervised embedding.

4 SBOW Comparison

Unlike Word2Vec, which only tries to preserve the similarity signal, the goals of DIVE cover preserving the capability of measuring not only the similarity but also whether one context distribution includes the other (inclusion signal) or being more general than the other (generality signal).

In this experiment, we perform a comprehensive comparison between SBOW and DIVE using multiple scoring functions to detect the hypernym relation between words based on different types of signal. The window size $|W|$ of SBOW is also set as 20, and experiment setups are the same as that described in Section 3.1. Notice that the comparison is inherently unfair because most of the information would be lost during the aggressive compression process of DIVE, and we would like to evaluate how well DIVE can preserve signals of interest using the number of dimensions which is several orders of magnitude less than that of SBOW.

4.1 Unsupervised Scoring Functions

After trying many existing and newly proposed functions which score a pair of words to detect hypernym relation between them, we find that good scoring functions for SBOW are also good scoring functions for DIVE. Thus, in addition to $C \cdot \Delta S$ used in Section 3.2, we also present 4 other best performing or representative scoring functions in the experiment (see our supplementary materials for more details):

³ Note that higher AP is reported for some models in previous literature: 80 (Vilnis and McCallum, 2015) in LEDS, 74.2 (Athiwaratkun and Wilson, 2017) in LEDS, and 20.6 (Vulić et al., 2016) in HyperLex. The difference could be caused by different train/test setup (e.g. How the hyper-parameters are tuned, different training corpus, etc.). However, DIVE beats even these results.

⁴<https://github.com/benathi/word2gm>

- **Inclusion:** CDE (Clarke, 2009) computes the summation of element-wise minimum over the magnitude of hyponym embedding (i.e. $\frac{\|\min(\mathbf{w}_p, \mathbf{w}_q)\|_1}{\|\mathbf{w}_q\|_1}$). CDE measures the degree of violation of equation (1). Equation (1) holds if and only if CDE is 1. Due to noise in SBOW, CDE is rarely exactly 1, but hypernym pairs usually have higher CDE. Despite its effectiveness, the good performance could mostly come from the magnitude of embeddings/features instead of inclusion properties among context distributions. To measure the inclusion properties between context distributions \mathbf{d}_p and \mathbf{d}_q (\mathbf{w}_p and \mathbf{w}_q after normalization, respectively), we use negative asymmetric L1 distance ($-AL_1$)⁵ as one of our scoring function, where

$$AL_1 = \min_a \sum_c w_0 \cdot \max(a\mathbf{d}_q[c] - \mathbf{d}_p[c], 0) + \max(\mathbf{d}_p[c] - a\mathbf{d}_q[c], 0), \quad (8)$$

and w_0 is a constant hyper-parameter.

- **Generality:** When the inclusion property in (2) holds, $\|\mathbf{y}\|_1 = \sum_i \mathbf{y}[i] \geq \sum_i \mathbf{x}[i] = \|\mathbf{x}\|_1$. Thus, we use summation difference ($\|\mathbf{w}_p\|_1 - \|\mathbf{w}_q\|_1$) as our score to measure generality signal (ΔS).
- **Similarity plus generality:** Computing cosine similarity on skip-grams (i.e. Word2Vec + C in Table 1) is a popular way to measure the similarity of two words, so we multiply the Word2Vec similarity with summation difference of DIVE or SBOW ($W \cdot \Delta S$) as an alternative of $C \cdot \Delta S$.

4.2 Baselines

- **SBOW Freq:** A word is represented by the frequency of its neighboring words. Applying PMI filter (set context feature to be 0 if its value is lower than $\log(k_f)$) to SBOW Freq only makes its performances closer to (but still much worse than) SBOW PPMI, so we omit the baseline.
- **SBOW PPMI:** SBOW which uses PPMI of its neighboring words as the features (Bullinaria and Levy, 2007). Applying PMI filter to SBOW PPMI usually makes the performances worse, especially when k_f is large. Similarly, a constant $\log(k')$ shifting to SBOW PPMI (i.e. $\max(PMI - \log(k'), 0)$) is not helpful, so we set both k_f and k' to be 1.

⁵The meaning and efficient implementation of AL_1 are illustrated in our supplementary materials

AP@all (%)		BLESS					EVALution					Lenci/Benotto				
		CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS
SBOW	Freq	6.3	7.3	5.6	11.0	5.9	35.3	32.6	36.2	33.0	36.3	51.8	47.6	51.0	51.8	51.1
	PPMI	13.6	5.1	5.6	17.2	15.3	30.4	27.7	34.1	31.9	34.3	47.2	39.7	50.8	51.1	52.0
	PPMI w/ IS	6.2	5.0	5.5	12.4	5.8	36.0	27.5	36.3	32.9	36.4	52.0	43.1	50.9	51.9	50.7
	All wiki	12.1	5.2	6.9	12.5	13.4	28.5	27.1	30.3	29.9	31.0	47.1	39.9	48.5	48.7	51.1
DIVE	Full	9.3	7.6	6.0	18.6	16.3	30.0	27.5	34.9	32.3	33.0	46.7	43.2	51.3	51.5	50.4
	w/o PMI	7.8	6.9	5.6	16.7	7.1	32.8	32.2	35.7	32.5	35.4	47.6	44.9	50.9	51.6	49.7
	w/o IS	9.0	6.2	7.3	6.2	7.3	24.3	25.0	22.9	23.5	23.9	38.8	38.1	38.2	38.2	38.4
Kmean (Freq NMF)		6.5	7.3	5.6	10.9	5.8	33.7	27.2	36.2	33.0	36.2	49.6	42.5	51.0	51.8	51.2

AP@all (%)		Weeds					Micro Average (4 datasets)					Medical				
		CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS
SBOW	Freq	69.5	58.0	68.8	68.2	68.4	23.1	21.8	22.9	25.0	23.0	19.4	19.2	14.1	18.4	15.3
	PPMI	61.0	50.3	70.3	69.2	69.3	24.7	17.9	22.3	28.1	27.8	23.4	8.7	13.2	20.1	24.4
	PPMI w/ IS	67.6	52.2	69.4	68.7	67.7	23.2	18.2	22.9	25.8	22.9	22.8	10.6	13.7	18.6	17.0
	All wiki	61.3	48.6	70.0	68.5	70.4	23.4	17.7	21.7	24.6	25.8	22.3	8.9	12.2	17.6	21.1
DIVE	Full	59.2	55.0	69.7	68.6	65.5	22.1	19.8	22.8	28.9	27.6	11.7	9.3	13.7	21.4	19.2
	w/o PMI	60.4	56.4	69.3	68.6	64.8	22.2	21.0	22.7	28.0	23.1	10.7	8.4	13.3	19.8	16.2
	w/o IS	49.2	47.3	45.1	45.1	44.9	18.9	17.3	17.2	16.8	17.5	10.9	9.8	7.4	7.6	7.7
Kmean (Freq NMF)		69.4	51.1	68.8	68.2	68.9	22.5	19.3	22.9	24.9	23.0	12.6	10.9	14.0	18.1	14.6

AP@all (%)		LEDS					TM14					Kotlerman 2010				
		CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS
SBOW	Freq	82.7	70.4	70.7	83.3	73.3	55.6	53.2	54.9	55.7	55.0	35.9	40.5	34.5	37.0	35.4
	PPMI	84.4	50.2	72.2	86.5	84.5	56.2	52.3	54.4	57.0	57.6	39.1	30.9	33.0	37.0	36.3
	PPMI w/ IS	81.6	54.5	71.0	84.7	73.1	57.1	51.5	55.1	56.2	55.4	37.4	31.0	34.4	37.8	35.9
	All wiki	83.1	49.7	67.9	82.9	81.4	54.7	50.5	52.6	55.1	54.9	38.5	31.2	32.2	35.4	35.3
DIVE	Full	83.3	74.7	72.7	86.4	83.5	55.3	52.6	55.2	57.3	57.2	35.3	31.6	33.6	37.4	36.6
	w/o PMI	79.3	74.8	72.0	85.5	78.7	54.7	53.9	54.9	56.5	55.4	35.4	38.9	33.8	37.8	36.7
	w/o IS	64.6	55.4	43.2	44.3	46.1	51.9	51.2	50.4	52.0	51.8	32.9	33.4	28.1	30.2	29.7
Kmean (Freq NMF)		80.3	64.5	70.7	83.0	73.0	54.8	49.0	54.8	55.6	54.8	32.1	37.0	34.5	36.9	34.8

AP@all (%)		HypeNet					WordNet					Micro Average (10 datasets)				
		CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS	CDE	AL ₁	ΔS	W·ΔS	C·ΔS
SBOW	Freq	37.5	28.3	46.9	35.9	43.4	56.6	55.2	55.5	56.2	55.6	31.1	28.2	31.5	31.6	31.2
	PPMI	23.8	24.0	47.0	32.5	33.1	57.7	53.9	55.6	56.8	57.2	30.1	23.0	31.1	32.9	33.5
	PPMI w/ IS	38.5	26.7	47.2	35.5	37.6	57.0	54.1	55.7	56.6	55.7	31.8	24.1	31.5	32.1	30.3
	All wiki	23.0	24.5	40.5	30.5	29.7	57.4	53.1	56.0	56.4	57.3	29.0	23.1	29.2	30.2	31.1
DIVE	Full	25.3	24.2	49.3	33.6	32.0	60.2	58.9	58.4	61.1	60.9	27.6	25.3	32.1	34.1	32.7
	w/o PMI	31.3	27.0	46.9	33.8	34.0	59.2	60.1	58.2	61.1	59.1	28.5	26.7	31.5	33.4	30.1
	w/o IS	20.1	21.7	20.3	21.8	22.0	61.0	56.3	51.3	55.7	54.7	22.3	20.7	19.1	19.6	19.9
Kmean (Freq NMF)		33.7	22.0	46.0	35.6	45.2	58.4	60.2	57.7	60.1	57.9	29.1	24.7	31.5	31.8	31.5

Table 2: AP@all (%) of 10 datasets. The box at lower right corner compares the micro average AP across all 10 datasets. Numbers in different rows come from different feature or embedding spaces. Numbers in different columns come from different datasets and unsupervised scoring functions. We also present the micro average AP across the first 4 datasets (BLESS, EVALution, Lenci/Benotto and Weeds), which are used as a benchmark for unsupervised hypernym detection (Shwartz et al., 2017). IS refers to inclusion shift on the shifted PMI matrix.

Spearman-ρ (%)		HyperLex				
		CDE	AL ₁	ΔS	W·ΔS	C·ΔS
SBOW	Freq	31.7	19.6	27.6	29.6	27.3
	PPMI	28.1	-2.3	31.8	34.3	34.5
	PPMI w/ IS	32.4	2.1	28.5	31.0	27.4
	All wiki	25.3	-2.2	28.0	30.5	31.0
DIVE	Full	28.9	18.7	31.2	33.3	32.8
	w/o PMI	29.2	22.2	29.5	31.9	29.2
	w/o IS	11.5	-0.9	-6.2	-10.0	-11.6
Kmean (Freq NMF)		30.6	3.3	27.5	29.5	27.6

Table 3: Spearman ρ (%) in HyperLex.

- SBOW PPMI w/ IS (with additional inclusion shift): The matrix reconstructed by DIVE when $k_I = 1$. Specifically, $\mathbf{w}[c] = \max(\log(\frac{P(w,c)}{P(w)*P(c)*\frac{Z}{\#(w)}}), 0)$.
- SBOW all wiki: SBOW using PPMI features trained on the whole WaCkypedia.
- DIVE without the PMI filter (DIVE w/o PMI)
- NMF on shifted PMI: Non-negative matrix factorization (NMF) on the shifted PMI without inclusion shift for DIVE (DIVE w/o IS). This is the same as applying the non-negative constraint on the skip-gram model.

SBOW Freq	SBOW PPMI	DIVE
5799	3808	20

Table 4: The average number of non-zero dimensions across all testing words in 10 datasets.

- K-means (Freq NMF): The method first uses Mini-batch k-means (Sculley, 2010) to cluster words in skip-gram embedding space into 100 topics, and hashes each frequency count in SBOW into the corresponding topic. If running k-means on skip-grams is viewed as an approximation of clustering the SBOW context vectors, the method can be viewed as a kind of NMF (Ding et al., 2005).

DIVE performs non-negative matrix factorization on PMI matrix after applying inclusion shift and PMI filtering. To demonstrate the effectiveness of each step, we show the performances of DIVE after removing PMI filtering (DIVE w/o PMI), removing inclusion shift (DIVE w/o IS), and removing matrix factorization (SBOW PPMI w/ IS, SBOW PPMI, and SBOW all wiki). The methods based on frequency matrix are also tested (SBOW Freq and Freq NMF).

4.3 Results and Discussions

In Table 2, we first confirm the finding of the previous review study of Shwartz et al. (2017): there is no single hypernymy scoring function which always outperforms others. One of the main reasons is that different datasets collect negative samples differently. For example, if negative samples come from random word pairs (e.g. WordNet dataset), a symmetric similarity measure is a good scoring function. On the other hand, negative samples come from related or similar words in HyperNet, EVALution, Lenci/Benotto, and Weeds, so only estimating generality difference leads to the best (or close to the best) performance. The negative samples in many datasets are composed of both random samples and similar words (such as BLESS), so the combination of similarity and generality difference yields the most stable results.

DIVE performs similar or better on most of the scoring functions compared with SBOW consistently across all datasets in Table 2 and Table 3, while using many fewer dimensions (see Table 4). This leads to 2-3 order of magnitude savings on both memory consumption and testing time. Furthermore, the low dimensional embedding makes the computational complexity independent of the vocabulary size, which drastically boosts the scalability of unsupervised hypernym detection especially with the help of GPU. It is surprising that we can achieve such aggressive compression while preserving the similarity, generality, and in-

clusion signal in various datasets with different types of negative samples. Its results on C- Δ S and W- Δ S outperform SBOW Freq. Meanwhile, its results on AL_1 outperform SBOW PPMI. The fact that W- Δ S or C- Δ S usually outperform generality functions suggests that only memorizing general words is not sufficient. The best average performance on 4 and 10 datasets are both produced by W- Δ S on DIVE.

SBOW PPMI improves the W- Δ S and C- Δ S from SBOW Freq but sacrifices AP on the inclusion functions. It generally hurts performance to directly include inclusion shift in PPMI (PPMI w/ IS) or compute SBOW PPMI on the whole Wikipedia (all wiki) instead of the first 51.2 million tokens. The similar trend can also be seen in Table 3. Note that AL_1 completely fails in the HyperLex dataset using SBOW PPMI, which suggests that PPMI might not necessarily preserve the distributional inclusion property, even though it can have good performance on scoring functions combining similarity and generality signals.

Removing the PMI filter from DIVE slightly drops the overall precision while removing inclusion shift on shifted PMI (w/o IS) leads to poor performances. K-means (Freq NMF) produces similar AP compared with SBOW Freq but has worse AL_1 scores. Its best AP scores on different datasets are also significantly worse than the best AP of DIVE. This means that only making Word2Vec (skip-grams) non-negative or naively accumulating topic distribution in contexts cannot lead to satisfactory embeddings.

5 Related Work

Most previous unsupervised approaches focus on designing better hypernymy scoring functions for sparse bag of word (SBOW) features. They are well summarized in the recent study (Shwartz et al., 2017). Shwartz et al. (2017) also evaluate the influence of different contexts, such as changing the window size of contexts or incorporating dependency parsing information, but neglect scalability issues inherent to SBOW methods.

A notable exception is the Gaussian embedding model (Vilnis and McCallum, 2015), which represents each word as a Gaussian distribution. However, since a Gaussian distribution is normalized, it is difficult to retain frequency information during the embedding process, and experiments on HyperLex (Vulić et al., 2016) demonstrate that a sim-

ple baseline only relying on word frequency can achieve good results. Follow-up work models contexts by a mixture of Gaussians (Athiwaratkun and Wilson, 2017) relaxing the unimodality assumption but achieves little improvement on hypernym detection tasks.

Kiela et al. (2015) show that images retrieved by a search engine can be a useful source of information to determine the generality of lexicons, but the resources (e.g. pre-trained image classifier for the words of interest) might not be available in many domains.

Order embedding (Vendrov et al., 2016) is a supervised approach to encode many annotated hypernym pairs (e.g. all of the whole WordNet (Miller, 1995)) into a compact embedding space, where the embedding of a hypernym should be smaller than the embedding of its hyponym in every dimension. Our method learns embedding from raw text, where a hypernym embedding should be larger than the embedding of its hyponym in every dimension. Thus, DIVE can be viewed as an unsupervised and reversed form of order embedding.

Non-negative matrix factorization (NMF) has a long history in NLP, for example in the construction of topic models (Pauca et al., 2004). Non-negative sparse embedding (NNSE) (Murphy et al., 2012) and Faruqui et al. (2015) indicate that non-negativity can make embeddings more interpretable and improve word similarity evaluations. The sparse NMF is also shown to be effective in cross-lingual lexical entailment tasks but does not necessarily improve monolingual hypernymy detection (Vyas and Carpuat, 2016). In our study, we show that performing NMF on PMI matrix with inclusion shift can preserve DIH in SBOW, and the comprehensive experimental analysis demonstrates its state-of-the-art performances on unsupervised hypernymy detection.

6 Conclusions

Although large SBOW vectors consistently show the best all-around performance in unsupervised hypernym detection, it is challenging to compress them into a compact representation which preserves inclusion, generality, and similarity signals for this task. Our experiments suggest that the existing approaches and simple baselines such as Gaussian embedding, accumulating K-mean clusters, and non-negative skip-grams do not lead to

satisfactory performance.

To achieve this goal, we propose an interpretable and scalable embedding method called *distributional inclusion vector embedding (DIVE)* by performing non-negative matrix factorization (NMF) on a weighted PMI matrix. We demonstrate that scoring functions which measure inclusion and generality properties in SBOW can also be applied to DIVE to detect hypernymy, and DIVE performs the best on average, slightly better than SBOW while using many fewer dimensions.

Our experiments also indicate that unsupervised scoring functions which combine similarity and generality measurements work the best in general, but no one scoring function dominates across all datasets. A combination of unsupervised DIVE with the proposed scoring functions produces new state-of-the-art performances on many datasets in the unsupervised regime.

7 Acknowledgement

This work was supported in part by the Center for Data Science and the Center for Intelligent Information Retrieval, in part by DARPA under agreement number FA8750-13-2-0020, in part by Defense Advanced Research Agency (DARPA) contract number HR0011-15-2-0036, in part by the National Science Foundation (NSF) grant numbers DMR-1534431 and IIS-1514053 and in part by the Chan Zuckerberg Initiative under the project Scientific Knowledge Base Construction. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, or the U.S. Government, or the other sponsors.

References

- Ben Athiwaratkun and Andrew Gordon Wilson. 2017. Multimodal word distributions. In *ACL*.
- Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In *EACL*.
- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. The WaCky wide web: a collection of very large linguistically processed

- web-crawled corpora. *Language resources and evaluation* 43(3):209–226.
- Marco Baroni and Alessandro Lenci. 2011. How we BLESSed distributional semantic evaluation. In *Workshop on GEometrical Models of Natural Language Semantics (GEMS)*.
- Giulia Benotto. 2015. Distributional models for semantic relations: A study on hyponymy and antonymy. *PhD Thesis, University of Pisa*.
- John A Bullinaria and Joseph P Levy. 2007. Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior research methods* 39(3):510–526.
- Philipp Cimiano, Andreas Hotho, and Steffen Staab. 2005. Learning concept hierarchies from text corpora using formal concept analysis. *J. Artif. Intell. Res.(JAIR)* 24(1):305–339.
- Daoud Clarke. 2009. Context-theoretic semantics for natural language: an overview. In *workshop on geometrical models of natural language semantics*. pages 112–119.
- Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. 2016. Lifted rule injection for relation embeddings. In *EMNLP*.
- Chris Ding, Xiaofeng He, and Horst D Simon. 2005. On the equivalence of nonnegative matrix factorization and spectral clustering. In *ICDM*.
- Manaal Faruqui, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah Smith. 2015. Sparse overcomplete word vector representations. In *ACL*.
- Edgar C Fieller, Herman O Hartley, and Egon S Pearson. 1957. Tests for rank correlation coefficients. i. *Biometrika*.
- Maayan Geffet and Ido Dagan. 2005. The distributional inclusion hypotheses and lexical entailment. In *ACL*.
- Goran Glavaš and Simone Paolo Ponzetto. 2017. Dual tensor model for detecting asymmetric lexico-semantic relations. In *EMNLP*.
- Zhiheng Huang, Marcus Thint, and Zengchang Qin. 2008. Question classification using head words and their hypernyms. In *EMNLP*.
- Douwe Kiela, Laura Rimell, Ivan Vulic, and Stephen Clark. 2015. Exploiting image generality for lexical entailment detection. In *ACL*.
- Diederik Kingma and Jimmy Ba. 2015. ADAM: A method for stochastic optimization. In *ICLR*.
- Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. 2010. Directional distributional similarity for lexical inference. *Natural Language Engineering* 16(4):359–389.
- Germán Kruszewski, Denis Paperno, and Marco Baroni. 2015. Deriving boolean structures from distributional vectors. *TACL* 3:375–388. <https://tacl2013.cs.columbia.edu/ojs/index.php/tacl/article/view/616>.
- Daniel D Lee and H Sebastian Seung. 2001. Algorithms for non-negative matrix factorization. In *NIPS*.
- Omer Levy, Ido Dagan, and Jacob Goldberger. 2014. Focused entailment graphs for open IE propositions. In *CoNLL*.
- Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *NIPS*.
- Omer Levy, Yoav Goldberg, and Ido Dagan. 2015a. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics* 3:211–225.
- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. 2015b. Do supervised distributional methods really learn lexical inference relations? In *NAACL-HTL*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- George A. Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM* 38(11):39–41.
- Brian Murphy, Partha Talukdar, and Tom Mitchell. 2012. Learning effective and interpretable semantic models using non-negative sparse embedding. *COLING* pages 1933–1950.
- Kim Anh Nguyen, Maximilian Köper, Sabine Schulte im Walde, and Ngoc Thang Vu. 2017. Hierarchical embeddings for hypernymy detection and directionality. In *EMNLP*.
- Maximilian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. In *NIPS*.
- V. Paul Pauca, Fariyal Shahnaz, Michael W Berry, and Robert J. Plemmons. 2004. Text mining using non-negative matrix factorizations. In *ICDM*.
- Boris Teodorovich Polyak. 1969. Minimization of unsmooth functionals. *USSR Computational Mathematics and Mathematical Physics*.
- Simone Paolo Ponzetto and Michael Strube. 2006. Exploiting semantic role labeling, wordnet and wikipedia for coreference resolution. In *ACL*.
- Stephen Roller and Katrin Erk. 2016. Relations such as hypernymy: Identifying and exploiting hearst patterns in distributional vectors for lexical entailment. In *EMNLP*.

- Mark Sammons, V Vydiswaran, and Dan Roth. 2011. Recognizing textual entailment. *Multilingual Natural Language Applications: From Theory to Practice*. Prentice Hall, Jun .
- Enrico Santus, Frances Yung, Alessandro Lenci, and Chu-Ren Huang. 2015. EVALution 1.0: an evolving semantic dataset for training and evaluation of distributional semantic models. In *Workshop on Linked Data in Linguistics (LDL)*.
- David Sculley. 2010. Web-scale k-means clustering. In *WWW*.
- Vered Shwartz, Yoav Goldberg, and Ido Dagan. 2016. Improving hypernymy detection with an integrated path-based and distributional method. In *ACL*.
- Vered Shwartz, Enrico Santus, and Dominik Schlechtweg. 2017. Hypernyms under siege: Linguistically-motivated artillery for hypernymy detection. In *EACL*.
- Peter D Turney and Saif M Mohammad. 2015. Experiments with three approaches to recognizing lexical entailment. *Natural Language Engineering* 21(3):437–476.
- Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2016. Order-embeddings of images and language. In *ICLR*.
- Luke Vilnis and Andrew McCallum. 2015. Word representations via gaussian embedding. In *ICLR*.
- Ivan Vulić, Daniela Gerz, Douwe Kiela, Felix Hill, and Anna Korhonen. 2016. Hyperlex: A large-scale evaluation of graded lexical entailment. *arXiv preprint arXiv:1608.02117* .
- Ivan Vulić and Nikola Mrkšić. 2017. Specialising word vectors for lexical entailment. *arXiv preprint arXiv:1710.06371* .
- Yogarshi Vyas and Marine Carpuat. 2016. Sparse bilingual word representations for cross-lingual lexical entailment. In *HLT-NAACL*.
- Julie Weeds, Daoud Clarke, Jeremy Reffin, David Weir, and Bill Keller. 2014. Learning to distinguish hypernyms and co-hyponyms. In *COLING*.
- Julie Weeds and David Weir. 2003. A general framework for distributional similarity. In *EMNLP*.
- Mu Zhu. 2004. Recall, precision and average precision. *Department of Statistics and Actuarial Science, University of Waterloo, Waterloo* 2:30.