

Learning Word Embeddings for Low-resource Languages by PU Learning

Chao Jiang

University of Virginia

cj7an@virginia.edu

Cho-Jui Hsieh

University of California Davis

chohsieh@ucdavis.edu

Hsiang-Fu Yu

Amazon

rofuyu@cs.utexas.edu

Kai-Wei Chang

University of California Los Angeles

kwchang@cs.ucla.edu

Abstract

Word embedding is a key component in many downstream applications in processing natural languages. Existing approaches often assume the existence of a large collection of text for learning effective word embedding. However, such a corpus may not be available for some low-resource languages. In this paper, we study how to effectively learn a word embedding model on a corpus with only a few million tokens. In such a situation, the co-occurrence matrix is sparse as the co-occurrences of many word pairs are unobserved. In contrast to existing approaches often only sample a few unobserved word pairs as negative samples, we argue that the zero entries in the co-occurrence matrix also provide valuable information. We then design a Positive-Unlabeled Learning (PU-Learning) approach to factorize the co-occurrence matrix and validate the proposed approaches in four different languages.

1 Introduction

Learning word representations has become a fundamental problem in processing natural languages. These semantic representations, which map a word into a point in a linear space, have been widely applied in downstream applications, including named entity recognition (Guo et al., 2014), document ranking (Nalisnick et al., 2016), sentiment analysis (Irsoy and Cardie, 2014), question answering (Antol et al., 2015), and image captioning (Karpathy and Fei-Fei, 2015).

Over the past few years, various approaches have been proposed to learn word vectors (e.g., (Pennington et al., 2014; Mikolov et al., 2013a; Levy and Goldberg, 2014b; Ji et al., 2015)) based on co-occurrence information between words observed on the training corpus. The intuition behind this is to represent words with similar vectors if

they have similar contexts. To learn a good word embedding, most approaches assume a large collection of text is freely available, such that the estimation of word co-occurrences is accurate. For example, the Google Word2Vec model (Mikolov et al., 2013a) is trained on the Google News dataset, which contains around 100 billion tokens, and the GloVe embedding (Pennington et al., 2014) is trained on a crawled corpus that contains 840 billion tokens in total. However, such an assumption may not hold for low-resource languages such as Inuit or Sindhi, which are not spoken by many people or have not been put into a digital format. For those languages, usually, only a limited size corpus is available. Training word vectors under such a setting is a challenging problem.

One key restriction of the existing approaches is that they often mainly rely on the word pairs that are observed to co-occur on the training data. When the size of the text corpus is small, most word pairs are unobserved, resulting in an extremely sparse co-occurrence matrix (i.e., most entries are zero)¹. For example, the text8² corpus has about 17,000,000 tokens and 71,000 distinct words. The corresponding co-occurrence matrix has more than five billion entries, but only about 45,000,000 are non-zeros (observed on the training corpus). Most existing approaches, such as Glove and Skip-gram, cannot handle a vast number of zero terms in the co-occurrence matrix; therefore, they only sub-sample a small subset of zero entries during the training.

In contrast, we argue that the *unobserved* word pairs can provide valuable information for training a word embedding model, especially when the co-occurrence matrix is very sparse. Inspired

¹Note that the zero term can mean either the pairs of words cannot co-occur or the co-occurrence is not observed in the training corpus.

²<http://mattmahoney.net/dc/text8.zip>

by the success of Positive-Unlabeled Learning (PU-Learning) in collaborative filtering applications (Pan et al., 2008; Hu et al., 2008; Pan and Scholz, 2009; Qin et al., 2010; Paquet and Koenigstein, 2013; Hsieh et al., 2015), we design an algorithm to effectively learn word embeddings from both positive (observed terms) and unlabeled (unobserved/zero terms) examples. Essentially, by using the square loss to model the unobserved terms and designing an efficient update rule based on linear algebra operations, the proposed PU-Learning framework can be trained efficiently and effectively.

We evaluate the performance of the proposed approach in English³ and other three resource-scarce languages. We collected unlabeled language corpora from Wikipedia and compared the proposed approach with popular approaches, the Glove and the Skip-gram models, for training word embeddings. The experimental results show that our approach significantly outperforms the baseline models, especially when the size of the training corpus is small.

Our key contributions are summarized below.

- We propose a PU-Learning framework for learning word embedding.
- We tailor the coordinate descent algorithm (Yu et al., 2017b) for solving the corresponding optimization problem.
- Our experimental results show that PU-Learning improves the word embedding training in the low-resource setting.

2 Related work

Learning word vectors. The idea of learning word representations can be traced back to Latent Semantic Analysis (LSA) (Deerwester et al., 1990) and Hyperspace Analogue to Language (HAL) (Lund and Burgess, 1996), where word vectors are generated by factorizing a word-document and word-word co-occurrence matrix, respectively. Similar approaches can also be extended to learn other types of relations between words (Yih et al., 2012; Chang et al., 2013) or entities (Chang et al., 2014). However, due to the limitation of the use of principal component analysis,

³Although English is not a resource-scarce language, we simulate the low-resource setting in an English corpus. In this way, we leverage the existing evaluation methods to evaluate the proposed approach.

these approaches are often less flexible. Besides, directly factorizing the co-occurrence matrix may cause the frequent words dominating the training objective.

In the past decade, various approaches have been proposed to improve the training of word embeddings. For example, instead of factorizing the co-occurrence count matrix, Bullinaria and Levy (2007); Levy and Goldberg (2014b) proposed to factorize *point-wise mutual information* (PMI) and *positive PMI* (PPMI) matrices as these metrics scale the co-occurrence counts (Bullinaria and Levy, 2007; Levy and Goldberg, 2014b). Skip-gram model with negative-sampling (SGNS) and Continuous Bag-of-Words models (Mikolov et al., 2013b) were proposed for training word vectors on a large scale without consuming a large amount of memory. GloVe (Pennington et al., 2014) is proposed as an alternative to decompose a weighted log co-occurrence matrix with a bias term added to each word. Very recently, *WordRank* model (Ji et al., 2015) has been proposed to minimize a ranking loss which naturally fits the tasks requiring ranking based evaluation metrics. Stratos et al. (2015) also proposed CCA (canonical correlation analysis)-based word embedding which shows competitive performance. All these approaches focus on the situations where a large text corpus is available.

Positive and Unlabeled (PU) Learning: Positive and Unlabeled (PU) learning (Li and Liu, 2005) is proposed for training a model when the positive instances are partially labeled and the unlabeled instances are mostly negative. Recently, PU learning has been used in many classification and collaborative filtering applications due to the nature of “implicit feedback” in many recommendation systems—users usually only provide positive feedback (e.g., purchases, clicks) and it is very hard to collect negative feedback.

To resolve this problem, a series of PU matrix completion algorithms have been proposed (Pan et al., 2008; Hu et al., 2008; Pan and Scholz, 2009; Qin et al., 2010; Paquet and Koenigstein, 2013; Hsieh et al., 2015; Yu et al., 2017b). The main idea is to assign a small uniform weight to all the missing or zero entries and factorize the corresponding matrix. Among them, Yu et al. (2017b) proposed an efficient algorithm for matrix factorization with PU-learning, such that the weighted matrix is constructed implicitly. In this paper, we

\mathcal{W}, \mathcal{C}	vocabulary of central and context words
m, n	vocabulary sizes
k	dimension of word vectors
W, H	$m \times k$ and $n \times k$ latent matrices
C_{ij}	weight for the (i, j) entry
A_{ij}	value of the PPMI matrix
Q_{ij}	value of the co-occurrence matrix
w_i, h_j	i -th row of W and j -th row of H
$\mathbf{b}, \hat{\mathbf{b}}$	bias term
$\lambda_i, \bar{\lambda}_j$	regularization parameters
$ \cdot $	the size of a set
Ω	Set of possible word-context pairs
Ω^+	Set of observed word-context pairs
Ω^-	Set of unobserved word-context pairs

Table 1: Notations.

design a new approach for training word vectors by leveraging the PU-Learning framework and existing word embedding techniques. To the best of our knowledge, this is the first work to train word embedding models using the PU-learning framework.

3 PU-Learning for Word Embedding

Similar to GloVe and other word embedding learning algorithms, the proposed approach consists of three steps. The first step is to construct a co-occurrence matrix. Follow the literature (Levy and Goldberg, 2014a), we use the PPMI metric to measure the co-occurrence between words. Then, in the second step, a PU-Learning approach is applied to factorize the co-occurrence matrix and generate word vectors and context vectors. Finally, a post-processing step generates the final embedding vector for each word by combining the word vector and the context vector.

We summarize the notations used in this paper in Table 1 and describe the details of each step in the remainder of this section.

3.1 Building the Co-Occurrence Matrix

Various metrics can be used for estimating the co-occurrence between words in a corpus. PPMI metric stems from *point-wise mutual information* (PMI) which has been widely used as a measure of word association in NLP for various tasks (Church and Hanks, 1990). In our case, each entry $PMI(w, c)$ represents the relevant measure between a word w and a context word c by calculating the ratio between their joint probability (the

chance they appear together in a local context window) and their marginal probabilities (the chance they appear independently) (Levy and Goldberg, 2014b). More specifically, each entry of PMI matrix can be defined by

$$PMI(w, c) = \log \frac{\hat{P}(w, c)}{\hat{P}(w) \cdot \hat{P}(c)}, \quad (1)$$

where $\hat{P}(w)$, $\hat{P}(c)$ and $\hat{P}(w, c)$ are the frequency of word w , word c , and word pairs (w, c) , respectively. The PMI matrix can be computed based on the co-occurrence counts of word pairs, and it is an information-theoretic association measure which effectively eliminates the big differences in magnitude among entries in the co-occurrence matrix.

Extending from the PMI metric, the PPMI metric replaces all the negative entries in PMI matrix by 0:

$$PPMI(w, c) = \max(PMI(w, c), 0). \quad (2)$$

The intuition behind this is that people usually perceive positive associations between words (e.g. “ice” and “snow”). In contrast, the negative association is hard to define (Levy and Goldberg, 2014b). Therefore, it is reasonable to replace the negative entries in the PMI matrix by 0, such that the negative association is treated as “uninformative”. Empirically, several existing works (Levy et al., 2015; Bullinaria and Levy, 2007) showed that the PPMI metric achieves good performance on various semantic similarity tasks.

In practice, we follow the pipeline described in Levy et al. (2015) to build the PPMI matrix and apply several useful tricks to improve its quality. First, we apply a context distribution smoothing mechanism to enlarge the probability of sampling a rare context. In particular, all context counts are scaled to the power of α .⁴:

$$PPMI_\alpha(w, c) = \max \left(\log \frac{\hat{P}(w, c)}{\hat{P}(w) \hat{P}_\alpha(c)}, 0 \right)$$

$$\hat{P}_\alpha(c) = \frac{\#(c)^\alpha}{\sum_{\bar{c}} \#(\bar{c})^\alpha},$$

where $\#(w)$ denotes the number of times word w appears. This smoothing mechanism effectively

⁴Empirically, $\alpha = 0.75$ works well (Mikolov et al., 2013b).

alleviates PPMI’s bias towards rare words (Levy et al., 2015).

Next, previous studies show that words that occur too frequent often dominate the training objective (Levy et al., 2015) and degrade the performance of word embedding. To avoid this issue, we follow Levy et al. (2015) to sub-sample words with frequency more than a threshold t with a probability p defined as:

$$p = 1 - \sqrt{\frac{t}{\hat{P}(w)}}.$$

3.2 PU-Learning for Matrix Factorization

We proposed a matrix factorization based word embedding model which aims to minimize the reconstruction error on the PPMI matrix. The low-rank embeddings are obtained by solving the following optimization problem:

$$\begin{aligned} \min_{W, H} \sum_{i, j \in \Omega} C_{ij} (A_{ij} - \mathbf{w}_i^T \mathbf{h}_j - b_i - \hat{b}_j)^2 \\ + \sum_i \lambda_i \|\mathbf{w}_i\|^2 + \sum_j \bar{\lambda}_j \|\mathbf{h}_j\|^2, \end{aligned} \quad (3)$$

where W and H are $m \times k$ and $n \times k$ latent matrices, representing words and context words, respectively. The first term in Eq. (3) aims for minimizing reconstruction error, and the second and third terms are regularization terms. λ_i and $\bar{\lambda}_j$ are weights of regularization term. They are hyperparameters that need to be tuned.

The zero entries in co-occurrence matrix denote that two words never appear together in the current corpus, which also refers to unobserved terms. The unobserved term can be either real zero (two words shouldn’t be co-occurred even when we use very large corpus) or just missing in the small corpus. In contrast to SGNS sub-sampling a small set of zero entries as negative samples, our model will try to use the information from all zeros.

The set Ω includes all the $|\mathcal{W}| \times |\mathcal{C}|$ entries—both positive and zero entries:

$$\Omega = \Omega^+ \cup \Omega^-. \quad (4)$$

Note that we define the positive samples Ω^+ to be all the (w, c) pairs that appear at least one time in the corpus, and negative samples Ω^- are word pairs that never appear in the corpus.

Weighting function. Eq (3) is very similar to the one used in previous matrix factorization approaches such as GloVe, but we propose a new way to set the weights C_{ij} . If we set equal weights for all the entries, then $C_{ij} = \text{constant}$, and the model is very similar to conducting SVD for the PPMI matrix. Previous work has shown that this approach often suffers from poor performance (Pennington et al., 2014). More advanced methods, such as GloVe, set non-uniform weights for observed entries to reflect their confidence. However, the time complexity of their algorithm is proportional to number of nonzero weights ($|(i, j) \mid C_{ij} \neq 0|$), thus they have to set zero weights for all the unobserved entries ($C_{ij} = 0$ for Ω^-), or try to incorporate a small set of unobserved entries by negative sampling.

We propose to set the weights for Ω^+ and Ω^- differently using the following scheme:

$$C_{ij} = \begin{cases} (Q_{ij}/x_{max})^\alpha, & \text{if } Q_{ij} \leq x_{max}, \text{ and } (i, j) \in \Omega^+ \\ 1, & \text{if } Q_{ij} > x_{max}, \text{ and } (i, j) \in \Omega^+ \\ \rho, & (i, j) \in \Omega^- \end{cases} \quad (5)$$

Here x_{max} and α are re-weighting parameters, and ρ is the unified weight for unobserved terms. We will discuss them later.

For entries in Ω^+ , we set the non-uniform weights as in GloVe (Pennington et al., 2014), which assigns larger weights to context word that appears more often with the given word, but also avoids overwhelming the other terms. For entries in Ω^- , instead of setting their weights to be 0, we assign a small constant weight ρ . The main idea is from the literature of PU-learning (Hu et al., 2008; Hsieh et al., 2015): although missing entries are highly uncertain, they are still likely to be true 0, so we should incorporate them in the learning process but multiplying with a smaller weight according to the uncertainty. Therefore, ρ in (5) reflects how confident we are to the zero entries.

In our experiments, we set $x_{max} = 10$, $\alpha = 3/4$ according to (Pennington et al., 2014), and let ρ be a parameter to tune. Experiments show that adding weighting function obviously improves the performance especially on analogy tasks.

Bias term. Unlike previous work on PU matrix completion (Yu et al., 2017b; Hsieh et al., 2015), we add the bias terms for word and context word

vectors. Instead of directly using $\mathbf{w}_i^\top \mathbf{h}_j$ to approximate A_{ij} , we use

$$A_{ij} \approx \mathbf{w}_i^\top \mathbf{h}_j + b_i + \hat{b}_j.$$

[Yu et al. \(2017b\)](#) design an efficient column-wise coordinate descent algorithm for solving the PU matrix factorization problem; however, they do not consider the bias term in their implementations. To incorporate the bias term in (3), we propose the following training algorithm based on the coordinate descent approach. Our algorithm does not introduce much overhead compared to that in [\(Yu et al., 2017b\)](#).

We augment each $\mathbf{w}_i, \mathbf{h}_j \in R^k$ into the following $(k+2)$ dimensional vectors:

$$\mathbf{w}'_i = \begin{bmatrix} w_{i1} \\ \vdots \\ w_{ik} \\ 1 \\ b_i \end{bmatrix} \quad \mathbf{h}'_j = \begin{bmatrix} h_{j1} \\ \vdots \\ h_{jk} \\ \hat{b}_j \\ 1 \end{bmatrix}$$

Therefore, for each word and context vector, we have the following equality

$$\langle \mathbf{w}'_i, \mathbf{h}'_j \rangle = \langle \mathbf{w}_i, \mathbf{h}_j \rangle + b_i + \hat{b}_j,$$

which means the loss function in (3) can be written as

$$\sum_{i,j \in \Omega} C_{ij} (A_{ij} - \mathbf{w}'_i^\top \mathbf{h}'_j)^2.$$

Also, we denote $W' = [\mathbf{w}'_1, \mathbf{w}'_2, \dots, \mathbf{w}'_n]^\top$ and $H' = [\mathbf{h}'_1, \mathbf{h}'_2, \dots, \mathbf{h}'_n]^\top$. In the column-wise coordinate descent method, at each iteration we pick a $t \in \{1, \dots, (k+2)\}$, and update the t -th column of W' and H' . The updates can be derived for the following two cases:

- When $t \leq k$, the elements in the t -th column is w_{1t}, \dots, w_{nt} and we can directly use the update rule derived in [Yu et al. \(2017b\)](#) to update them.
- When $t = k+1$, we do not update the corresponding column of W' since the elements are all 1, and we use the similar coordinate descent update to update the $k+1$ -th column of H' (corresponding to $\hat{b}_1, \dots, \hat{b}_n$). When $t = k+2$, we do not update the corresponding column of H' (they are all 1) and we update the $k+2$ -th column of W' (corresponding to b_1, \dots, b_n) using coordinate descent.

With some further derivations, we can show that the algorithm only requires $O(\text{nnz}(A) + nk)$ time to update each column,⁵ so the overall complexity is $O(\text{nnz}(A)k + nk^2)$ time per epoch, which is only proportional to number of nonzero terms in A . Therefore, with the same time complexity as GloVe, we can utilize the information from all the zero entries in A instead of only sub-sampling a small set of zero entries.

3.3 Interpretation of Parameters

In the PU-Learning formulation, ρ represents the unified weight that assigned to the unobserved terms. Intuitively, ρ reflects the confidence on unobserved entries—larger ρ means that we are quite certain about the zeroes, while small ρ indicates the many of unobserved pairs are not truly zero. When $\rho = 0$, the PU-Learning approach reduces to a model similar to GloVe, which discards all the unobserved terms. In practice, ρ is an important parameter to tune, and we find that $\rho = 0.0625$ achieves the best results in general. Regarding the other parameter, λ is the regularization term for preventing the embedding model from overfitting. In practice, we found the performance is not very sensitive to λ as long as it is reasonably small. More discussion about the parameter setting can be found in Section 5.

Post-processing of Word/Context Vectors The PU-Learning framework factorizes the PPMI matrix and generates two vectors for each word i , $\mathbf{w}_i \in R^k$ and $\mathbf{h}_i \in R^k$. The former represents the word when it is the central word and the latter represents the word when it is in context. [Levy et al. \(2015\)](#) shows that averaging these two vectors ($\mathbf{u}_i^{\text{avg}} = \mathbf{w}_i + \mathbf{h}_i$) leads to consistently better performance. The same trick of constructing word vectors is also used in GloVe. Therefore, in the experiments, we evaluate all models with \mathbf{u}^{avg} .

4 Experimental Setup

Our goal in this paper is to train word embedding models for low-resource languages. In this section, we describe the experimental designs to evaluate the proposed PU-learning approach. We first describe the data sets and the evaluation metrics. Then, we provide details of parameter tuning.

⁵Here we assume $m = n$ for the sake of simplicity. And, $\text{nnz}(A)$ denotes the number of nonzero terms in the matrix A .

	Similarity task					Analogy task	
Word embedding	WS353	Similarity	Relatedness	M. Turk	MEN	3CosAdd	3CosMul
GloVe	48.7	50.9	53.7	54.1	17.6	32.1	28.5
SGNS	67.2	70.3	67.9	59.9*	25.1*	30.4	27.8
PU-learning	68.3*	71.8*	68.2*	57.0	22.7	32.6*	32.3*

Table 2: Performance of the best SGNS, GloVe, PU-Learning models, trained on the text8 corpus. Results show that our proposed model is better than SGNS and GloVe. Star indicates it is significantly better than the second best algorithm in the same column according to Wilcoxon signed-rank test. ($p < 0.05$)

Language	Similarity task					Analogy task
	WS353	Similarity	Relatedness	M. Turk	MEN	
English (en)	353	203	252	287	3,000	19,544
Czech (cs)	337	193	241	268	2,810	18,650
Danish (da)	346	198	247	283	2,951	18,340
Dutch (nl)	346	200	247	279	2,852	17,684

Table 3: The size of the test sets. The data sets in English are the original test sets. To evaluate other languages, we translate the data sets from English.

4.1 Evaluation tasks

We consider two widely used tasks for evaluating word embeddings, the word similarity task and the word analogy task. In the word similarity task, each question contains a word pairs and an annotated similarity score. The goal is to predict the similarity score between two words based on the inner product between the corresponding word vectors. The performance is then measured by the Spearman’s rank correlation coefficient, which estimates the correlation between the model predictions and human annotations. Following the settings in literature, the experiments are conducted on five data sets, *WordSim353* (Finkelstein et al., 2001), *WordSim Similarity* (Zesch et al., 2008), *WordSim Relatedness* (Agirre et al., 2009), *Mechanical Turk* (Radinsky et al., 2011) and *MEN* (Bruni et al., 2012).

In the word analogy task, we aim at solving analogy puzzles like “man is to woman as king is to ?”, where the expected answer is “queen.” We consider two approaches for generating answers to the puzzles, namely 3CosAdd and 3CosMul (see (Levy and Goldberg, 2014a) for details). We evaluate the performances on *Google analogy* dataset (Mikolov et al., 2013a) which contains 8,860 semantic and 10,675 syntactic questions. For the analogy task, only the answer that exactly matches the annotated answer is counted as correct. As a result, the analogy task is more difficult than the similarity task because the eval-

uation metric is stricter and it requires algorithms to differentiate words with similar meaning and find the right answer.

To evaluate the performances of models in the low-resource setting, we train word embedding models on Dutch, Danish, Czech and, English data sets collected from Wikipedia. The original Wikipedia corpora in Dutch, Danish, Czech and English contain 216 million, 47 million, 92 million, and 1.8 billion tokens, respectively. To simulate the low-resource setting, we sub-sample the Wikipedia corpora and create a subset of 64 million tokens for Dutch and Czech and a subset of 32 million tokens for English. We will demonstrate how the size of the corpus affects the performance of embedding models in the experiments.

To evaluate the performance of word embeddings in Czech, Danish, and Dutch, we translate the English similarity and analogy test sets to the other languages by using Google Cloud Translation API⁶. However, an English word may be translated to multiple words in another language (e.g., compound nouns). We discard questions containing such words (see Table 3 for details). Because all approaches are compared on the same test set for each language, the comparisons are fair.

4.2 Implementation and Parameter Setting

We compare the proposed approach with two baseline methods, GloVe and SGNS. The imple-

⁶<https://cloud.google.com/translate>

Dutch (nl)	Similarity task					Analogy task	
Word embedding	WS353	Similarity	Relatedness	M. Turk	MEN	3CosAdd	3CosMul
GloVe	35.4	35.0	41.7	44.3	11	21.2	20.2
SGNS	51.9	52.9	53.5	49.8*	15.4	22.1	23.6
PU-learning	53.7*	53.4*	55.1*	46.7	16.4*	23.5*	24.7*
Danish (da)	Similarity task					Analogy task	
Word embedding	WS353	Similarity	Relatedness	M. Turk	MEN	3CosAdd	3CosMul
GloVe	25.7	18.4	40.3	49.0	16.4	25.8*	24.3*
SGNS	49.7	47.1	52.1	51.5	22.4	22.0	21.2
PU-learning	53.5*	49.5*	59.3*	51.7*	22.7*	22.6	22.8
Czech (cs)	Similarity task					Analogy task	
Word embedding	WS353	Similarity	Relatedness	M. Turk	MEN	3CosAdd	3CosMul
GloVe	34.3	23.2	48.9	36.5	16.2	8.9	8.6
SGNS	51.4	42.7	61.1	44.2	21.3	10.4*	9.8
PU-learning	54.0*	45.4*	65.3*	46.2*	21.7*	9.9	10.1*
English (en)	Similarity task					Analogy task	
Word embedding	WS353	Similarity	Relatedness	M. Turk	MEN	3CosAdd	3CosMul
GloVe	47.9	52.1	49.5	58.8	19.1	34.3	32.6
SGNS	65.7	67.1*	66.5	62.8*	26.1*	31.2	27.4
PU-learning	67.0*	66.7	69.6*	59.4	22.4	39.2*	38.8*

Table 4: Performance of SGNS, GloVe, and the proposed PU-Learning model in four different languages. Results show that the proposed PU-Learning model outperforms SGNS and GloVe in most cases when the size of corpus is relatively small (around 50 million tokens). Star indicates it is significant better than the second best algorithm in the same column according to Wilcoxon signed-rank test. ($p < 0.05$).

mentations of Glove⁷ and SGNS⁸ and provided by the original authors, and we apply the default settings when appropriate. The proposed PU-Learning framework is implemented based on Yu et al. (2017a). With the implementation of efficient update rules, our model requires less than 500 seconds to perform one iteration over the entire text8 corpus, which consists of 17 million tokens⁹. All the models are implemented in C++.

We follow Levy et al. (2015)¹⁰ to set windows size as 15, minimal count as 5, and dimension of word vectors as 300 in the experiments. Training word embedding models involves selecting several hyper-parameters. However, as the word embeddings are usually evaluated in an unsupervised setting (i.e., the evaluation data sets are not seen during the training), the parameters should not be tuned on each dataset. To conduct a fair comparison, we tune hyper-parameters on the text8 dataset. For GloVe model, we tune the discount parameters x_{max} and find that $x_{max} = 10$ per-

forms the best. SGNS has a natural parameter k which denotes the number of negative samples. Same as Levy et al. (2015), we found that setting k to 5 leads to the best performance. For the PU-learning model, ρ and λ are two important parameters that denote the unified weight of zero entries and the weight of regularization terms, respectively. We tune ρ in a range from 2^{-1} to 2^{-14} and λ in a range from 2^0 to 2^{-10} . We analyze the sensitivity of the model to these hyper-parameters in the experimental result section. The best performance of each model on the text8 dataset is shown in the Table 2. It shows that PU-learning model outperforms two baseline models.

5 Experimental Results

We compared the proposed PU-Learning framework with two popular word embedding models – SGNS (Mikolov et al., 2013b) and Glove (Pennington et al., 2014) on English and three other languages. The experimental results are reported in Table 4. The results show that the proposed PU-Learning framework outperforms the two baseline approaches significantly in most datasets. This re-

⁷<https://nlp.stanford.edu/projects/glove>

⁸<https://code.google.com/archive/p/word2vec/>

⁹<http://mattmahoney.net/dc/text8.zip>

¹⁰<https://bitbucket.org/omerlevy/hyperwords>

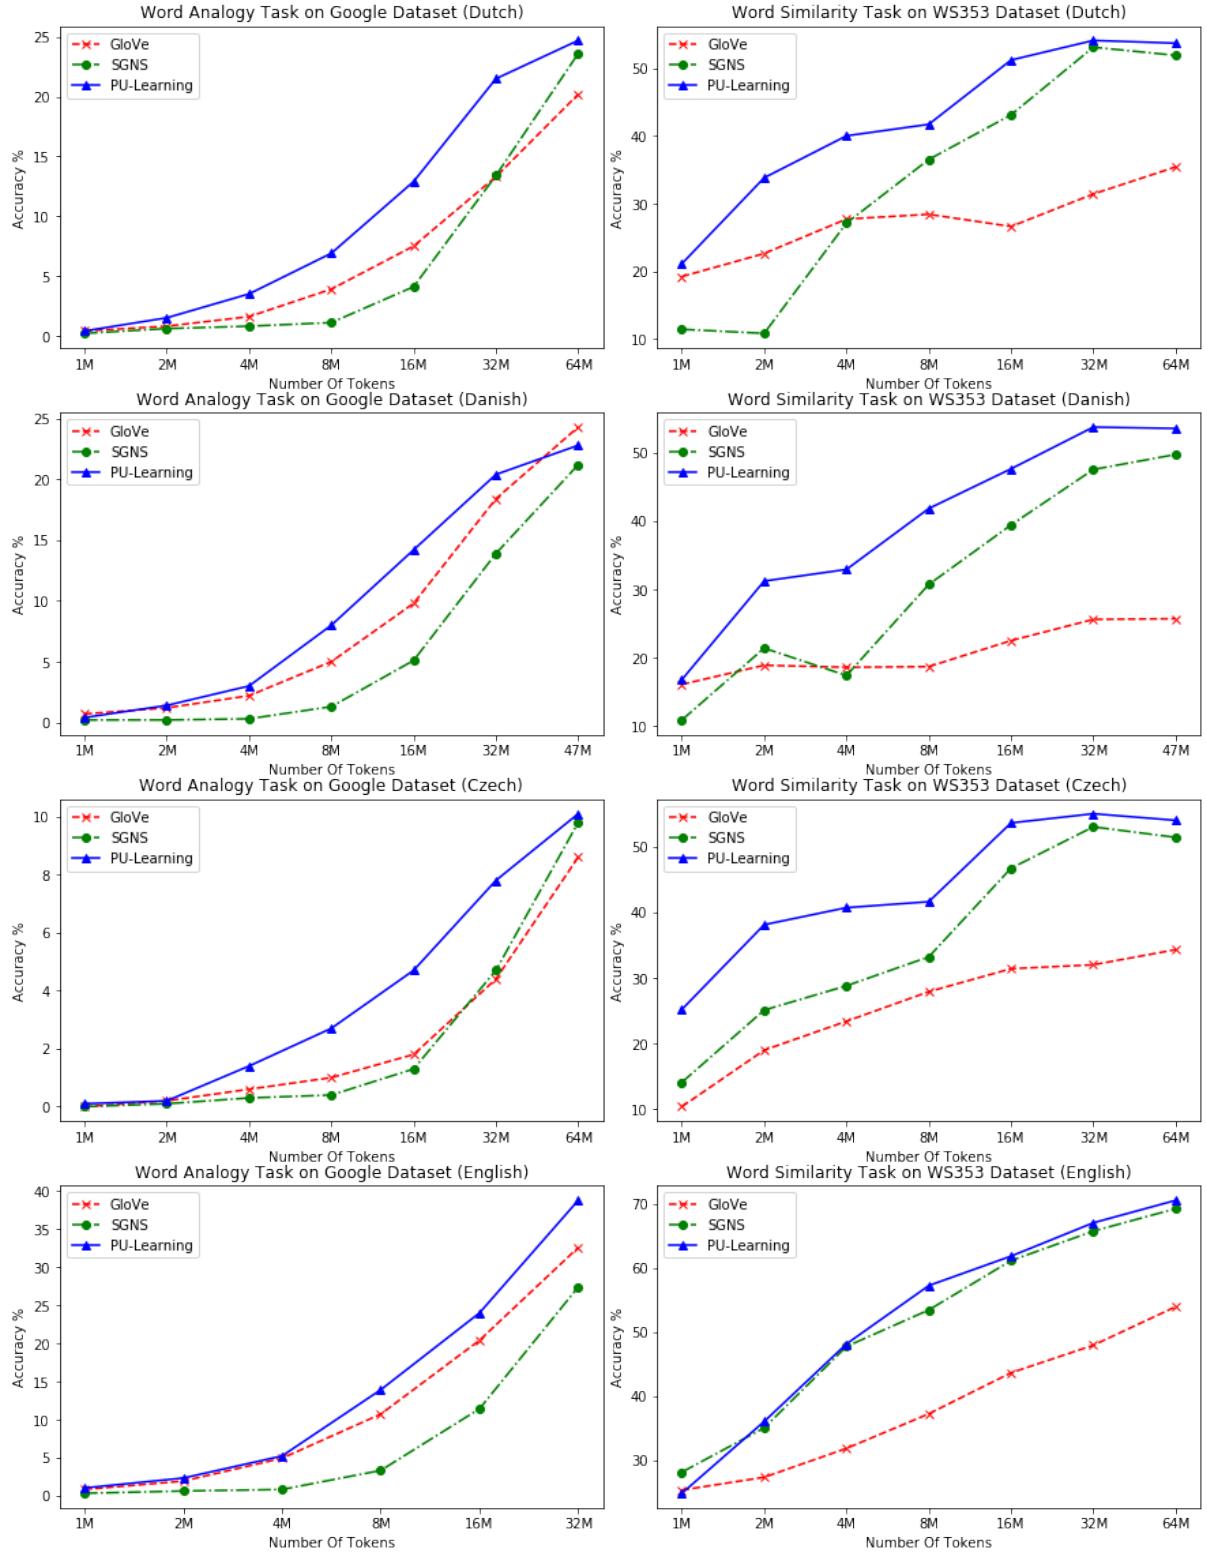


Figure 1: Performance change as the corpus size growing (a) on the Google word analogy task (on the left-hand side) and (b) on the WS353 word similarity task (on the right-hand side). We demonstrate the performance on four languages, Dutch, Danish, Czech and English datasets. Results show that PU-Learning model consistently outperforms SGNS and GloVe when the size of corpus is small.

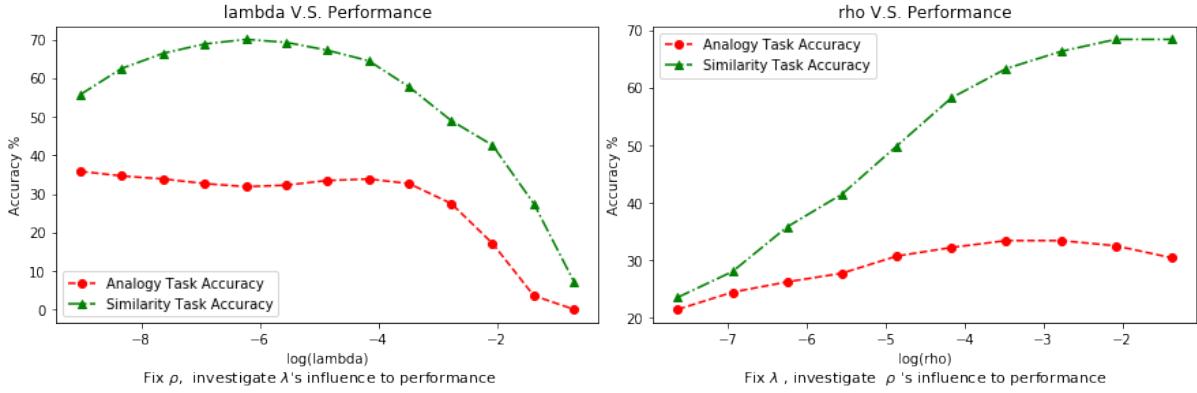


Figure 2: Impact of ρ and λ in the PU-Learning framework.

sults confirm that the unobserved word pairs carry important information and the PU-Learning model leverages such information and achieves better performance. To better understand the model, we conduct detailed analysis as follows.

Performance v.s. Corpus size We investigate the performance of our algorithm with respect to different corpus size, and plot the results in Figure 1. The results in analogy task are obtained by 3CosMul method (Levy and Goldberg, 2014a). As the corpus size grows, the performance of all models improves, and the PU-learning model consistently outperforms other methods in all the tasks. However, with the size of the corpus increases, the difference becomes smaller. This is reasonable as when the corpus size increases the number of non-zero terms becomes smaller and the PU-learning approach is resemblance to Glove.

Impacts of ρ and λ We investigate how sensitive the model is to the hyper-parameters, ρ and λ . Figure 2 shows the performance along with various values of λ and ρ when training on the text8 corpus, respectively. Note that the x-axis is in log scale. When ρ is fixed, a big λ degrades the performance of the model significantly. This is because when λ is too big the model suffers from underfitting. The model is less sensitive when λ is small and in general, $\lambda = 2^{-11}$ achieves consistently good performance.

When λ is fixed, we observe that large ρ (e.g., $\rho \approx 2^{-4}$) leads to better performance. As ρ represents the weight assigned to the unobserved term, this result confirms that the model benefits from using the zero terms in the co-occurrences matrix.

6 Conclusion

In this paper, we presented a PU-Learning framework for learning word embeddings of low-resource languages. We evaluated the proposed approach on English and other three languages and showed that the proposed approach outperforms other baselines by effectively leveraging the information from unobserved word pairs.

In the future, we would like to conduct experiments on other languages where available text corpora are relatively hard to obtain. We are also interested in applying the proposed approach to domains, such as legal documents and clinical notes, where the amount of accessible data is small. Besides, we plan to study how to leverage other information to facilitate the training of word embeddings under the low-resource setting.

Acknowledge

This work was supported in part by National Science Foundation Grant IIS-1760523, IIS-1719097 and an NVIDIA Hardware Grant.

References

Eneko Agirre, Enrique Alfonseca, Keith Hall, Jana Kravalova, Marius Paşa, and Aitor Soroa. 2009. A study on similarity and relatedness using distributional and WordNet-based approaches. In *The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. pages 19–27.

Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual Question Answering. In *Proceedings of the IEEE International Conference on Computer Vision*. pages 2425–2433.

Elia Bruni, Gemma Boleda, Marco Baroni, and Nam-Khanh Tran. 2012. Distributional semantics in technicolor. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*. Association for Computational Linguistics, pages 136–145.

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.

Kai-Wei Chang, Wen tau Yih, Bishan Yang, and Chris Meek. 2014. Typed Tensor Decomposition of Knowledge Bases for Relation Extraction. In *EMNLP*.

Kai-Wei Chang, Wen-tau Yih, and Christopher Meek. 2013. Multi-relational latent semantic analysis. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. pages 1602–1612.

Kenneth Ward Church and Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational linguistics* 16(1):22–29.

Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American society for information science* 41(6):391.

Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*. ACM, pages 406–414.

Jiang Guo, Wanxiang Che, Haifeng Wang, and Ting Liu. 2014. Revisiting embedding features for simple semi-supervised learning. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pages 110–120.

Cho-Jui Hsieh, Nagarajan Natarajan, and Inderjit Dhillon. 2015. Pu learning for matrix completion. In *International Conference on Machine Learning*. pages 2445–2453.

Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative filtering for implicit feedback datasets. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*. pages 263–272.

Ozan Irsoy and Claire Cardie. 2014. Deep recursive neural networks for compositionality in language. In *Advances in neural information processing systems*. pages 2096–2104.

Shihao Ji, Hyokun Yun, Pinar Yanardag, Shin Matsushima, and SVN Vishwanathan. 2015. Wordinrank: Learning word embeddings via robust ranking. *arXiv preprint arXiv:1506.02761* .

Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 3128–3137.

Omer Levy and Yoav Goldberg. 2014a. Linguistic regularities in sparse and explicit word representations. In *Proceedings of the eighteenth conference on computational natural language learning*. pages 171–180.

Omer Levy and Yoav Goldberg. 2014b. Neural word embedding as implicit matrix factorization. In *Advances in Neural Information Processing Systems*. pages 2177–2185.

Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics* 3:211–225.

Xiao-Li Li and Bing Liu. 2005. Learning from positive and unlabeled examples with different data distributions. In *European Conference on Machine Learning*. Springer, pages 218–229.

Kevin Lund and Curt Burgess. 1996. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers* 28(2):203–208.

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.

Eric Nalisnick, Bhaskar Mitra, Nick Craswell, and Rich Caruana. 2016. Improving document ranking with dual word embeddings. In *Proceedings of the 25th International Conference Companion on World Wide Web*. International World Wide Web Conferences Steering Committee, pages 83–84.

Rong Pan and Martin Scholz. 2009. Mind the gaps: Weighting the unknown in large-scale one-class collaborative filtering. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. pages 667–676.

Rong Pan, Yunhong Zhou, Bin Cao, Nathan N Liu, Rajan Lukose, Martin Scholz, and Qiang Yang. 2008. One-class collaborative filtering. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on*. IEEE, pages 502–511.

Ulrich Paquet and Noam Koenigstein. 2013. One-class collaborative filtering with random graphs. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, pages 999–1008.

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

Tao Qin, Tie-Yan Liu, Jun Xu, and Hang Li. 2010. LETOR: A benchmark collection for research on learning to rank for information retrieval. *Information Retrieval* 13(4):346–374.

Kira Radinsky, Eugene Agichtein, Evgeniy Gabrilovich, and Shaul Markovitch. 2011. A word at a time: computing word relatedness using temporal semantic analysis. In *Proceedings of the 20th international conference on World wide web*. ACM, pages 337–346.

Karl Stratos, Michael Collins, and Daniel Hsu. 2015. Model-based word embeddings from decompositions of count matrices. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1282–1291.

Wen-tau Yih, Geoffrey Zweig, and John C Platt. 2012. Polarity inducing latent semantic analysis. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. Association for Computational Linguistics, pages 1212–1222.

Hsiang-Fu Yu, Mikhail Bilenko, and Chih-Jen Lin. 2017a. Selection of negative samples for one-class matrix factorization. In *Proceedings of the 2017 SIAM International Conference on Data Mining*. SIAM, pages 363–371.

Hsiang-Fu Yu, Hsin-Yuan Huang, Inderjit S Dhillon, and Chih-Jen Lin. 2017b. A unified algorithm for one-class structured matrix factorization with side information. In *AAAI*, pages 2845–2851.

Torsten Zesch, Christof Müller, and Iryna Gurevych. 2008. Using wiktionary for computing semantic relatedness. In *AAAI*, volume 8, pages 861–866.