

Cross-lingual Models of Word Embeddings: An Empirical Comparison

Shyam Upadhyay¹ Manaal Faruqui² Chris Dyer² Dan Roth¹

¹ Department of Computer Science, University of Illinois, Urbana-Champaign, IL, USA

² School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

upadhya3@illinois.edu, mfaruqui@cs.cmu.edu

cdyer@cs.cmu.edu, danr@illinois.edu

Abstract

Despite interest in using cross-lingual knowledge to learn word embeddings for various tasks, a systematic comparison of the possible approaches is lacking in the literature. We perform an extensive evaluation of four popular approaches of inducing cross-lingual embeddings, each requiring a different form of supervision, on four typologically different language pairs. Our evaluation setup spans four different tasks, including intrinsic evaluation on mono-lingual and cross-lingual similarity, and extrinsic evaluation on downstream semantic and syntactic applications. We show that models which require expensive cross-lingual knowledge almost always perform better, but cheaply supervised models often prove competitive on certain tasks.

1 Introduction

Learning word vector representations using mono-lingual distributional information is now a ubiquitous technique in NLP. The quality of these word vectors can be significantly improved by incorporating cross-lingual distributional information (Klementiev et al., 2012; Zou et al., 2013; Vulić and Moens, 2013b; Mikolov et al., 2013b; Faruqui and Dyer, 2014; Hermann and Blunsom, 2014; Chandar et al., 2014, *inter alia*), with improvements observed both on monolingual (Faruqui and Dyer, 2014; Rastogi et al., 2015) and cross-lingual tasks (Guo et al., 2015; Søgaard et al., 2015; Guo et al., 2016).


Several models for inducing cross-lingual embeddings have been proposed, each requiring a different form of cross-lingual supervision – some can use document-level alignments (Vulić and Moens, 2015), others need alignments at the sentence (Hermann and Blunsom, 2014; Gouws et

al., 2015) or word level (Faruqui and Dyer, 2014; Gouws and Søgaard, 2015), while some require both sentence and word alignments (Luong et al., 2015). However, a systematic comparison of these models is missing from the literature, making it difficult to analyze which approach is suitable for a particular NLP task. In this paper, we fill this void by empirically comparing four cross-lingual word embedding models each of which require different form of alignment(s) as supervision, across several dimensions. To this end, we train these models on four different language pairs, and evaluate them on both monolingual and cross-lingual tasks.¹

First, we show that different models can be viewed as instances of a more general framework for inducing cross-lingual word embeddings. Then, we evaluate these models on both extrinsic and intrinsic tasks. Our intrinsic evaluation assesses the quality of the vectors on monolingual (§4.2) and cross-lingual (§4.3) word similarity tasks, while our extrinsic evaluation spans semantic (cross-lingual document classification §4.4) and syntactic tasks (cross-lingual dependency parsing §4.5).

Our experiments show that word vectors trained using expensive cross-lingual supervision (word alignments or sentence alignments) perform the best on semantic tasks. On the other hand, for syntactic tasks like cross-lingual dependency parsing, models requiring weaker form of cross-lingual supervision (such as context agnostic translation dictionary) are competitive to models requiring expensive supervision. We also show qualitatively how the nature of cross-lingual supervision used to train word vectors affects the proximity of translation pairs across languages, and of words with similar meaning in the same language in the vector-space.

¹Instructions and code to reproduce the experiments available at http://cogcomp.cs.illinois.edu/page/publication_view/794



Algorithm 1 General Algorithm

- 1: Initialize $\mathbf{W} \leftarrow \mathbf{W}^0, \mathbf{V} \leftarrow \mathbf{V}^0$
 - 2: $(\mathbf{W}^*, \mathbf{V}^*) \leftarrow \arg \min \alpha A(\mathbf{W}) + \beta B(\mathbf{V}) + C(\mathbf{W}, \mathbf{V})$
-

Figure 1: **(Above)** A general schema for induction of cross-lingual word vector representations. The word vector model generates embeddings which incorporates distributional information *cross-lingually*. **(Below)** A general algorithm for inducing bilingual word embeddings, where $\alpha, \beta, \mathbf{W}^0, \mathbf{V}^0$ are parameters and A, B, C are suitably defined losses.

2 Bilingual Embeddings

A general schema for inducing bilingual embeddings is shown in Figure 1. Our comparison focuses on dense, fixed-length distributed embeddings which are obtained using some form of cross-lingual supervision. We briefly describe the embedding induction procedure for each of the selected bilingual word vector models, with the aim to provide a unified algorithmic perspective for all methods, and to facilitate better understanding and comparison. Our choice of models spans across different forms of supervision required for inducing the embeddings, illustrated in Figure 2.

Notation. Let $W = \{w_1, w_2, \dots, w_{|W|}\}$ be the vocabulary of a language l_1 with $|W|$ words, and $\mathbf{W} \in \mathbb{R}^{|W| \times l}$ be the corresponding word embeddings of length l . Let $V = \{v_1, v_2, \dots, v_{|V|}\}$ be the vocabulary of another language l_2 with $|V|$ words, and $\mathbf{V} \in \mathbb{R}^{|V| \times m}$ the corresponding word embeddings of length m . We denote the word vector for a word w by \mathbf{w} .

2.1 Bilingual Skip-Gram Model (BiSkip)

Luong et al. (2015) proposed Bilingual Skip-Gram, a simple extension of the monolingual skip-gram model, which learns bilingual embeddings by using a parallel corpus along with word alignments (both sentence and word level alignments).

The learning objective is a simple extension of the skip-gram model, where the context of a word is expanded to include bilingual links obtained from word alignments, so that the model is trained to predict words cross-lingually. In par-

ticular, given a word alignment link from word $v \in V$ in language l_2 to $w \in W$ in language l_1 , the model predicts the context words of w using v and vice-versa. Formally, the cross lingual part of the objective is,

$$D_{12}(\mathbf{W}, \mathbf{V}) = - \sum_{(v,w) \in Q} \sum_{w_c \in \text{NBR}_1(w)} \log P(w_c | v) \quad (1)$$

where $\text{NBR}_1(w)$ is the context of w in language l_1 , Q is the set of word alignments, and $P(w_c | v) \propto \exp(\mathbf{w}_c^T \mathbf{v})$. Another similar term D_{21} models the objective for v and $\text{NBR}_2(v)$. The objective can be cast into Algorithm 1 as,

$$C(\mathbf{W}, \mathbf{V}) = D_{12}(\mathbf{W}, \mathbf{V}) + D_{21}(\mathbf{W}, \mathbf{V}) \quad (2)$$

$$A(\mathbf{W}) = - \sum_{w \in W} \sum_{w_c \in \text{NBR}_1(w)} \log P(w_c | w) \quad (3)$$

$$B(\mathbf{V}) = - \sum_{v \in V} \sum_{v_c \in \text{NBR}_2(v)} \log P(v_c | v) \quad (4)$$

where $A(\mathbf{W})$ and $B(\mathbf{V})$ are the familiar skip-gram formulation of the monolingual part of the objective. α and β are chosen hyper-parameters which set the relative importance of the monolingual terms.

2.2 Bilingual Compositional Model (BiCVM)

Hermann and Blunsom (2014) present a method that learns bilingual word vectors from a sentence aligned corpus. Their model leverages the fact that aligned sentences have equivalent meaning, thus their sentence representations should be similar.

We denote two aligned sentences, $\vec{v} = \langle \mathbf{x}_1, \dots \rangle$ and $\vec{w} = \langle \mathbf{y}_1, \dots \rangle$, where $\mathbf{x}_i \in \mathbf{V}, \mathbf{y}_i \in \mathbf{W}$, are vectors corresponding to the words in the sentences. Let functions $f : \vec{v} \rightarrow \mathbb{R}^n$ and $g : \vec{w} \rightarrow \mathbb{R}^n$, map sentences to their semantic representations in \mathbb{R}^n . BiCVM generates word vectors by minimizing the squared ℓ_2 norm between the sentence representations of aligned sentences. In order to prevent the degeneracy arising from directly minimizing the ℓ_2 norm, they use a noise-contrastive large-margin update, with randomly drawn sentence pairs (\vec{v}, \vec{w}^n) as negative samples. The loss for the sentence pairs (\vec{v}, \vec{w}) and (\vec{v}, \vec{w}^n) can be written as,

$$E(\vec{v}, \vec{w}, \vec{w}^n) = \max(\delta + \Delta E(\vec{v}, \vec{w}, \vec{w}^n), 0) \quad (5)$$

where,

$$E(\vec{v}, \vec{w}) = \|f(\vec{v}) - g(\vec{w})\|^2 \quad (6)$$




Figure 2: Forms of supervision required by the four models compared in this paper. From left to right, the cost of the supervision required varies from expensive (BiSkip) to cheap (BiVCD). BiSkip requires a parallel corpus annotated with word alignments (Fig. 2a), BiCVM requires a sentence-aligned corpus (Fig. 2b), BiCCA only requires a bilingual lexicon (Fig. 2c) and BiVCD requires comparable documents (Fig. 2d).

and,

$$\Delta E(\vec{v}, \vec{w}, \vec{w}^n) = E(\vec{v}, \vec{w}) - E(\vec{v}, \vec{w}^n) \quad (7)$$

This can be cast into Algorithm 1 by,

$$C(\mathbf{W}, \mathbf{V}) = \sum_{\substack{\text{aligned } (\vec{v}, \vec{w}) \\ \text{random } \vec{w}^n}} E(\vec{v}, \vec{w}, \vec{w}^n) \quad (8)$$

$$A(\mathbf{W}) = \|\mathbf{W}\|^2 \quad B(\mathbf{V}) = \|\mathbf{V}\|^2 \quad (9)$$

with $A(\mathbf{W})$ and $B(\mathbf{V})$ being regularizers, with $\alpha = \beta$.

2.3 Bilingual Correlation Based Embeddings (BiCCA)

The BiCCA model, proposed by Faruqui and Dyer (2014), showed that when (independently trained) monolingual vector matrices \mathbf{W}, \mathbf{V} are projected using CCA (Hotelling, 1936) to respect a translation lexicon, their performance improves on word similarity and word analogy tasks. They first construct $\mathbf{W}' \subseteq \mathbf{W}, \mathbf{V}' \subseteq \mathbf{V}$ such that $|\mathbf{W}'| = |\mathbf{V}'|$ and the corresponding words (w_i, v_i) in the matrices are translations of each other. The projection is then computed as:

$$\mathbf{P}_W, \mathbf{P}_V = \text{CCA}(\mathbf{W}', \mathbf{V}') \quad (10)$$

$$\mathbf{W}^* = \mathbf{W}\mathbf{P}_W \quad \mathbf{V}^* = \mathbf{V}\mathbf{P}_V \quad (11)$$

where, $\mathbf{P}_V \in \mathbb{R}^{l \times d}, \mathbf{P}_W \in \mathbb{R}^{m \times d}$ are the projection matrices with $d \leq \min(l, m)$ and the $\mathbf{V}^* \in \mathbb{R}^{|\mathbf{V}'| \times d}, \mathbf{W}^* \in \mathbb{R}^{|\mathbf{W}'| \times d}$ are the word vectors that have been “enriched” using bilingual knowledge.

The BiCCA objective can be viewed² as the following instantiation of Algorithm 1:

$$\mathbf{W}^0 = \mathbf{W}', \mathbf{V}^0 = \mathbf{V}' \quad (12)$$

$$C(\mathbf{W}, \mathbf{V}) = \|\mathbf{W} - \mathbf{V}\|^2 + \gamma (\mathbf{V}^T \mathbf{W}) \quad (13)$$

$$A(\mathbf{W}) = \|\mathbf{W}\|^2 - 1 \quad B(\mathbf{V}) = \|\mathbf{V}\|^2 - 1 \quad (14)$$

where $\mathbf{W} = \mathbf{W}^0 \mathbf{P}_W$ and $\mathbf{V} = \mathbf{V}^0 \mathbf{P}_V$, where we set $\alpha = \beta = \gamma = \infty$ to set hard constraints.

²described in Section 6.5 of (Hardoon et al., 2004)

2.4 Bilingual Vectors from Comparable Data (BiVCD)

Another approach of inducing bilingual word vectors, which we refer to as BiVCD, was proposed by Vulić and Moens (2015). Their approach is designed to use *comparable* corpus between the source and target language pair to induce cross-lingual vectors.

Let d_e and d_f denote a pair of comparable documents with length in words p and q respectively (assume $p > q$). BiVCD first merges these two comparable documents into a single pseudo-bilingual document using a deterministic strategy based on length ratio of two documents $R = \lfloor \frac{p}{q} \rfloor$. Every R^{th} word of the merged pseudo-bilingual document is picked sequentially from d_f . Finally, a skip-gram model is trained on the corpus of pseudo-bilingual documents, to generate vectors for all words in $\mathbf{W}^* \cup \mathbf{V}^*$. The vectors constituting \mathbf{W}^* and \mathbf{V}^* can then be easily identified.

Instantiating BiVCD in the general algorithm is obvious: $C(\mathbf{W}, \mathbf{V})$ assumes the familiar word2vec skip-gram objective over the pseudo-bilingual document,

$$C(\mathbf{W}, \mathbf{V}) = - \sum_{s \in W \cup V} \sum_{t \in \text{NBR}(s)} \log P(t | s) \quad (15)$$

where $\text{NBR}(s)$ is defined by the pseudo-bilingual document and $P(t | s) \propto \exp(\mathbf{t}^T \mathbf{s})$. Note that $t, s \in W \cup V$.

Although BiVCD is designed to use comparable corpus, we provide it with parallel data in our experiments (to ensure comparability), and treat two aligned sentences as comparable.

3 Data

We train cross-lingual embeddings for 4 language pairs: English-German (en-de), English-French (en-fr), English-Swedish (en-sv) and English-Chinese (en-zh). For en-de and en-sv we use the

l_1	l_2	#sent	# l_1 -words	# l_2 -words
	de	1.9	53	51
en	fr	2.0	55	61
	sv	1.7	46	42
	zh	2.0	58	50

Table 1: The size of parallel corpora (in millions) of different language pairs used for training cross-lingual word vectors.

Europarl v7 parallel corpus³ (Koehn, 2005). For en-fr, we use Europarl combined with the news-commentary and UN-corpus dataset from WMT 2015.⁴ For en-zh, we use the FBIS parallel corpus from the news domain (LDC2003E14). We use the Stanford Chinese Segmenter (Tseng et al., 2005) to preprocess the en-zh parallel corpus. Corpus statistics for all languages is shown in Table 1.

4 Evaluation

We measure the quality of the induced cross-lingual word embeddings in terms of their performance, when used as features in the following tasks:

- monolingual word similarity for English
- Cross-lingual dictionary induction
- Cross-lingual document classification
- Cross-lingual syntactic dependency parsing

The first two tasks intrinsically measure how much can monolingual and cross-lingual similarity benefit from cross-lingual training. The last two tasks measure the ability of cross-lingually trained vectors to extrinsically facilitate model transfer across languages, for semantic and syntactic applications respectively. These tasks have been used in previous works (Klementiev et al., 2012; Luong et al., 2015; Vulić and Moens, 2013a; Guo et al., 2015) for evaluating cross-lingual embeddings, but no comparison exists which uses them in conjunction.

To ensure fair comparison, all models are trained with embeddings of size 200. We provide all models with parallel corpora, irrespective of their requirements. Whenever possible, we also report statistical significance of our results.

³www.statmt.org/europarl/v7/{de,sv}-en.tgz

⁴www.statmt.org/wmt15/translation-task.html

4.1 Parameter Selection

We follow the **BestAvg** parameter selection strategy from Lu et al. (2015): we selected the parameters for all models by tuning on a set of values (described below) and picking the parameter setting which did best on an average across all tasks.

BiSkip. All models were trained using a window size of 10 (tuned over $\{5, 10, 20\}$), and 30 negative samples (tuned over $\{10, 20, 30\}$). The cross-lingual weight was set to 4 (tuned over $\{1, 2, 4, 8\}$). The word alignments for training the model (available at github.com/lmthang/bivec) were generated using `fast_align` (Dyer et al., 2013). The number of training iterations was set to 5 (no tuning) and we set $\alpha = 1$ and $\beta = 1$ (no tuning).

BiCVM. We use the tool (available at github.com/karlmoritz/bicvm) released by Hermann and Blunsom (2014) to train all embeddings. We train an additive model (that is, $f(\vec{x}) = g(\vec{x}) = \sum_i x_i$) with hinge loss margin set to 200 (no tuning), batch size of 50 (tuned over 50, 100, 1000) and noise parameter of 10 (tuned over $\{10, 20, 30\}$). All models are trained for 100 iterations (no tuning).

BiCCA. First, monolingual word vectors are trained using the skip-gram model⁵ with negative sampling (Mikolov et al., 2013a) with window of size 5 (tuned over $\{5, 10, 20\}$). To generate a cross-lingual dictionary, word alignments are generated using `cdec` from the parallel corpus. Then, word pairs (a, b) , $a \in l_1, b \in l_2$ are selected such that a is aligned to b the most number of times and vice versa. This way, we obtained dictionaries of approximately 36k, 35k, 30k and 28k word pairs for en-de, en-fr, en-sv and en-zh respectively.

The monolingual vectors are aligned using the above dictionaries with the tool (available at github.com/mfaruqui/eacl14-cca) released by Faruqui and Dyer (2014) to generate the cross-lingual word embeddings. We use $k = 0.5$ as the number of canonical components (tuned over $\{0.2, 0.3, 0.5, 1.0\}$). Note that this results in a embedding of size 100 after performing CCA.

BiVCD. We use `word2vec`'s skip gram model for training our embeddings, with a window size of 5 (tuned on $\{5, 10, 20, 30\}$) and negative sampling parameter set to 5 (tuned on $\{5, 10, 25\}$). Every pair of parallel sentences is treated as a

⁵code.google.com/p/word2vec

pair of comparable documents, and merging is performed using the sentence length ratio strategy described earlier.⁶

4.2 Monolingual Evaluation

We first evaluate if the inclusion of cross-lingual knowledge improves the quality of English embeddings.

Word Similarity. Word similarity datasets contain word pairs which are assigned similarity ratings by humans. The task evaluates how well the notion of word similarity according to humans is emulated in the vector space. Evaluation is based on the Spearman’s rank correlation coefficient (Myers and Well, 1995) between human rankings and rankings produced by computing cosine similarity between the vectors of two words.

We use the SimLex dataset for English (Hill et al., 2014) which contains 999 pairs of English words, with a balanced set of noun, adjective and verb pairs. SimLex is claimed to capture word similarity exclusively instead of WordSim-353 (Finkelstein et al., 2001) which captures both word similarity and relatedness. We declare significant improvement if $p < 0.1$ according to Steiger’s method (Steiger, 1980) for calculating the statistical significant differences between two dependent correlation coefficients.

Table 2 shows the performance of English embeddings induced by all the models by training on different language pairs on the SimLex word similarity task. The score obtained by monolingual English embeddings trained on the respective English side of each language is shown in column marked Mono. In all cases (except BiCCA on en-sv), the bilingually trained vectors achieve better scores than the mono-lingually trained vectors.

Overall, across all language pairs, BiCVM is the best performing model in terms of Spearman’s correlation, but its improvement over BiSkip and BiVCD is often insignificant. It is notable that 2 of the 3 top performing models, BiCVM and BiVCD, need sentence aligned and document-aligned corpus only, which are easier to obtain than parallel data with word alignments required by BiSkip.

QVEC. Tsvetkov et al. (2015) proposed an intrinsic evaluation metric for estimating the quality of English word vectors. The score produced by QVEC measures how well a given set of word vectors is able to quantify linguistic properties

⁶We implemented the code for performing the merging as we could not find a tool provided by the authors.

pair	Mono	BiSkip	BiCVM	BiCCA	BiVCD
en-de	0.29	<u>0.34</u>	0.37	0.30	0.32
en-fr	0.30	<u>0.35</u>	0.39	0.31	0.36
en-sv	0.28	<u>0.32</u>	0.34	0.27	<u>0.32</u>
en-zh	0.28	<u>0.34</u>	0.39	0.30	0.31
avg.	0.29	0.34	0.37	0.30	0.33

Table 2: Word similarity score measured in Spearman’s correlation ratio for English on SimLex-999. The best score for each language pair is shown in **bold**. Scores which are significantly better (per Steiger’s Method with $p < 0.1$) than the next lower score are underlined. For example, for en-zh, BiCVM is significantly better than BiSkip, which in turn is significantly better than BiVCD.

pair	Mono	BiSkip	BiCVM	BiCCA	BiVCD
en-de	0.39	0.40	0.31	0.33	0.37
en-fr	0.39	0.40	0.31	0.33	0.38
en-sv	0.39	0.39	0.31	0.32	0.37
en-zh	0.40	0.40	0.32	0.33	0.38
avg.	0.39	0.40	0.31	0.33	0.38

Table 3: Intrinsic evaluation of English word vectors measured in terms of QVEC score across models. Best scores for each language pair is shown in **bold**.

of words, with higher being better. The metric is shown to have strong correlation with performance on downstream semantic applications. As it can be currently only used for English, we use it to evaluate the English vectors obtained using cross-lingual training of different models. Table 3 shows that on average across language pairs, BiSkip achieves the best score, followed by Mono (mono-lingually trained English vectors), BiVCD and BiCCA. A possible explanation for why Mono scores are better than those obtained by some of the cross-lingual models is that QVEC measures monolingual semantic content based on a linguistic oracle made for English. Cross-lingual training might affect these semantic properties arbitrarily.

Interestingly, BiCVM which was the best model according to SimLex, ranks last according to QVEC. The fact that the best models according to QVEC and word similarities are different reinforces observations made in previous work that performance on word similarity tasks alone does not reflect quantification of linguistic properties of words (Tsvetkov et al., 2015; Schnabel et al., 2015).

l_1	l_2	BiSkip	BiCVM	BiCCA	BiVCD
en	de	79.7	74.5	72.4	62.5
	fr	78.9	72.9	70.1	68.8
	sv	77.1	76.7	74.2	56.9
	zh	69.4	66.0	59.6	53.2
avg.	76.3	72.5	69.1	60.4	

Table 4: Cross-lingual dictionary induction results (top-10 accuracy). The same trend was also observed across models when computing MRR (mean reciprocal rank).

4.3 Cross-lingual Dictionary Induction

The task of cross-lingual dictionary induction (Vulić and Moens, 2013a; Gouws et al., 2015; Mikolov et al., 2013b) judges how good cross-lingual embeddings are at detecting word pairs that are semantically similar across languages. We follow the setup of Vulić and Moens (2013a), but instead of manually creating a gold cross-lingual dictionary, we derived our gold dictionaries using the Open Multilingual WordNet data released by Bond and Foster (2013). The data includes synset alignments across 26 languages with over 90% accuracy. First, we prune out words from each synset whose frequency count is less than 1000 in the vocabulary of the training data from §3. Then, for each pair of aligned synsets $s_1 = \{k_1, k_2, \dots\}$ $s_2 = \{g_1, g_2, \dots\}$, we include all elements from the set $\{(k, g) \mid k \in s_1, g \in s_2\}$ into the gold dictionary, where k and g are the lemmas. Using this approach we generated dictionaries of sizes 1.5k, 1.4k, 1.0k and 1.6k pairs for en-fr, en-de, en-sv and en-zh respectively.

We report top-10 accuracy, which is the fraction of the entries (e, f) in the gold dictionary, for which f belongs to the list of top-10 neighbors of the word vector of e , according to the induced cross-lingual embeddings. From the results (Table 4), it can be seen that for dictionary induction, the performance improves with the quality of supervision. As we move from cheaply supervised methods (eg. BiVCD) to more expensive supervision (eg. BiSkip), the accuracy improves. This suggests that for cross lingual similarity tasks, the more expensive the cross-lingual knowledge available, the better. Models using weak supervision like BiVCD perform poorly in comparison to models like BiSkip and BiCVM, with performance gaps upwards of 10 pts on an average.

l_1	l_2	BiSkip	BiCVM	BiCCA	BiVCD
en	de	85.2	<u>85.0</u>	79.1	79.9
	fr	77.7	71.7	70.7	72.0
	sv	72.3	<u>69.1</u>	<u>65.3</u>	59.9
	zh	75.5	73.6	69.4	<u>73.0</u>
de		74.9	<u>71.1</u>	64.9	<u>74.1</u>
fr	en	80.4	73.7	<u>75.5</u>	<u>77.6</u>
sv		<u>73.4</u>	67.7	67.0	78.2
zh		81.1	76.4	77.3	<u>80.9</u>
avg.		77.6	73.5	71.2	74.5

Table 5: Cross-lingual document classification accuracy when trained on language l_1 , and evaluated on language l_2 . The best score for each language is shown in **bold**. Scores which are significantly better (per McNemar’s Test with $p < 0.05$) than the next lower score are underlined. For example, for $sv \rightarrow en$, BiVCD is significantly better than BiSkip, which in turn is significantly better than BiCVM.

4.4 Cross-lingual Document Classification

We follow the cross-lingual document classification (CLDC) setup of Klementiev et al. (2012), but extend it to cover all of our language pairs. We use the RCV2 Reuters multilingual corpus⁷ for our experiments. In this task, for a language pair (l_1, l_2) , a document classifier is trained using the document representations derived from word embeddings in language l_1 , and then the trained model is tested on documents from language l_2 (and vice-versa). By using supervised training data in one language and evaluating without further supervision in another, CLDC assesses whether the learned cross-lingual representations are semantically coherent across multiple languages.

All embeddings are learned on the data described in §3, and we only use the RCV2 data to learn document classification models. Following previous work, we compute document representation by taking the tf-idf weighted average of vectors of the words present in it.⁸ A multi-class classifier is trained using an averaged perceptron (Freund and Schapire, 1999) for 10 iterations, using the document vectors of language l_1 as features⁹. Majority baselines for $en \rightarrow l_2$ and $l_1 \rightarrow en$ are 49.7% and 46.7% respectively, for all languages. Table 5 shows the performance of different models across different language pairs. We computed confidence values using the McNemar test (McNe-

⁷<http://trec.nist.gov/data/reuters/reuters.html>

⁸tf-idf (Salton and Buckley, 1988) was computed using all documents for that language in RCV2.

⁹We use the implementation of Klementiev et al. (2012).

mar, 1947) and declare significant improvement if $p < 0.05$.

Table 5 shows that in almost all cases, BiSkip performs significantly better than the remaining models. For transferring semantic knowledge across languages via embeddings, sentence and word level alignment proves superior to sentence or word level alignment alone. This observation is consistent with the trend in cross-lingual dictionary induction, where too the most expensive form of supervision performed the best.

4.5 Cross-lingual Dependency Parsing

Using cross lingual similarity for direct-transfer of dependency parsers was first shown in Täckström et al. (2012). The idea behind direct-transfer is to train a dependency parsing model using embeddings for language l_1 and then test the trained model on language l_2 , replacing embeddings for language l_1 with those of l_2 . The transfer relies on coherence of the embeddings across languages arising from the cross lingual training. For our experiments, we use the cross lingual transfer setup of Guo et al. (2015).¹⁰ Their framework trains a transition-based dependency parser using nonlinear activation function, with the source-side embeddings as lexical features. These embeddings can be replaced by target-side embeddings at test time.

All models are trained for 5000 iterations with fixed word embeddings during training. Since our goal is to determine the utility of word embeddings in dependency parsing, we turn off other features that can capture distributional information like brown clusters, which were originally used in Guo et al. (2015). We use the universal dependency treebank (McDonald et al., 2013) version-2.0 for our evaluation. For Chinese, we use the treebank released as part of the CoNLL-X shared task (Buchholz and Marsi, 2006).

We first evaluate how useful the word embeddings are in cross-lingual model transfer of dependency parsers (Table 6). On an average, BiCCA does better than other models. BiSkip is a close second, with an average performance gap of less than 1 point. BiSkip outperforms BiCVM on German and French (over 2 point improvement), owing to word alignment information BiSkip’s model uses during training. It is not surprising that English-Chinese transfer scores are low, due to the significant difference in syntactic structure of the

¹⁰github.com/jiangfeng1124/acl15-clnndep

l_1	l_2	BiSkip	BiCVM	BiCCA	BiVCD
en	de	49.8	47.5	51.3	49.0
	fr	65.8	63.2	65.9	60.7
	sv	56.9	56.7	59.4	54.6
	zh	6.4	6.1	6.4	6.0
de		49.7	45.0	50.3	43.6
fr		53.3	50.6	54.2	49.5
sv	en	48.2	49.0	49.9	44.6
zh		0.17	0.12	0.17	0.15
avg.		41.3	39.8	42.2	38.5

Table 6: Labeled attachment score (LAS) for cross-lingual dependency parsing when trained on language l_1 , and evaluated on language l_2 . The best score for each language is shown in **bold**.

two languages. Surprisingly, unlike the semantic tasks considered earlier, the models with expensive supervision requirements like BiSkip and BiCVM could not outperform a cheaply supervised BiCCA.

We also evaluate whether using cross-lingually trained vectors for learning dependency parsers is better than using mono-lingually trained vectors in Table 7. We compare against parsing models trained using mono-lingually trained word vectors (column marked Mono in Table 7). These vectors are the same used as input to the BiCCA model. All other settings remain the same. On an average across language pairs, improvement over the monolingual embeddings was obtained with the BiSkip and BiCCA models, while BiCVM and BiVCD consistently performed worse. A possible reason for this is that BiCVM and BiVCD operate on sentence level contexts to learn the embeddings, which only captures the semantic meaning of the sentences and ignores the internal syntactic structure. As a result, embedding trained using BiCVM and BiVCD are not informative for syntactic tasks. On the other hand, BiSkip and BiCCA both utilize the word alignment information to train their embeddings and thus do better in capturing some notion of syntax.

5 Qualitative Analysis

Figure 3 shows the PCA projection of some of the most frequent words in the English-French corpus. It is clear that BiSkip and BiCVM produce cross-lingual vectors which are the most comparable, the English and French words which are translations of each other are represented by almost the same point in the vector-space. In BiCCA and BiVCD

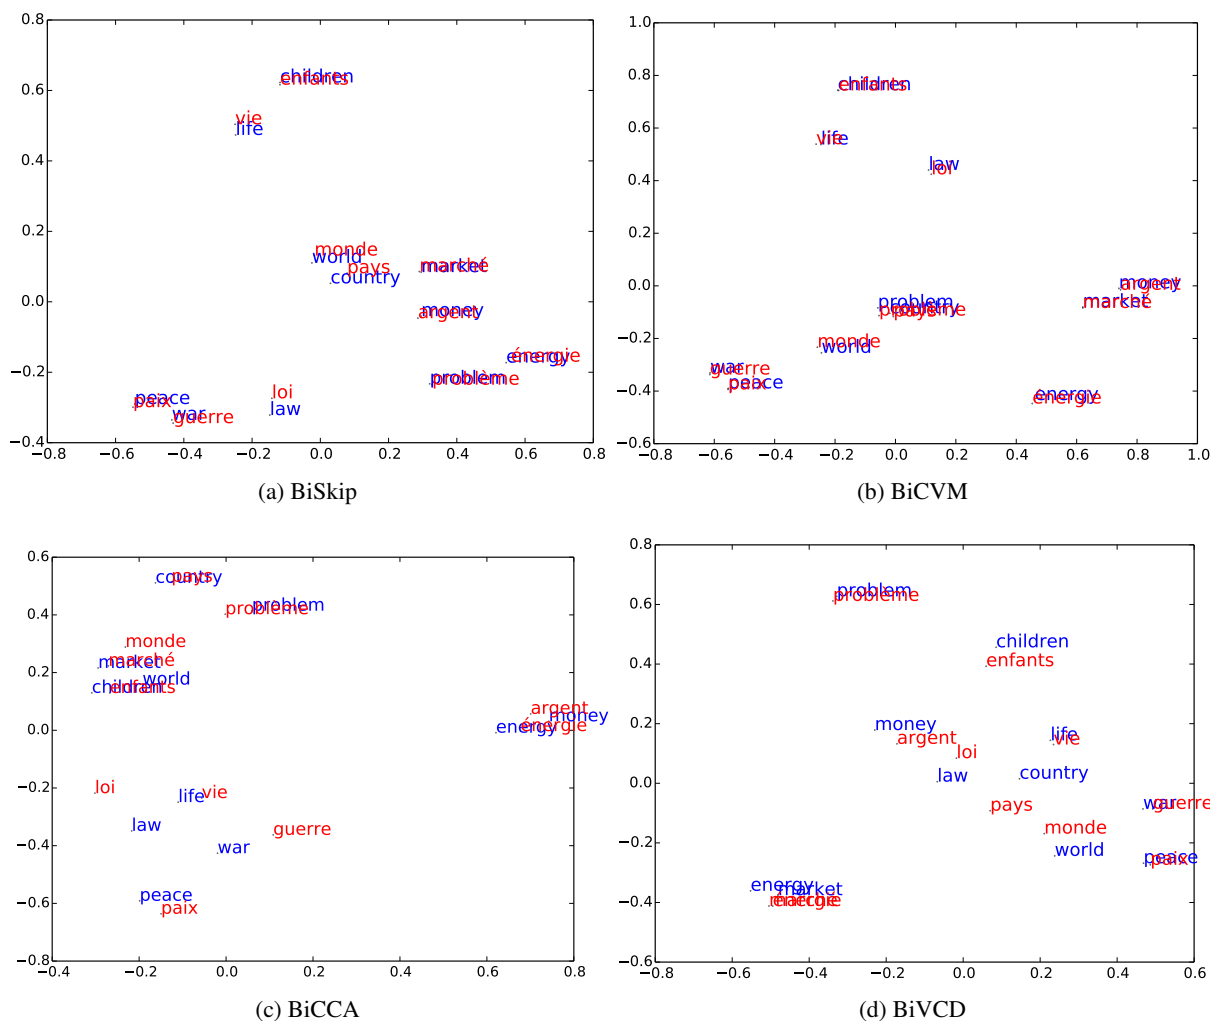


Figure 3: PCA projection of word embeddings of some frequent words present in English-French corpus. English and French words are shown in blue and red respectively.

l	Mono	BiSkip	BiCVM	BiCCA	BiVCD
de	71.1	72.0	60.4	71.4	58.9
fr	78.9	80.4	73.7	80.2	69.5
sv	75.5	78.2	70.5	79.0	64.5
zh	73.8	73.1	65.8	71.7	67.0
avg.	74.8	75.9	67.6	75.6	66.8

Table 7: Labeled attachment score (LAS) for dependency parsing when trained and tested on language l . Mono refers to parser trained with mono-lingually induced embeddings. Scores in **bold** are better than the Mono scores for each language, showing improvement from cross-lingual training.

the translated words are more distant than BiSkip and BiCVM. This is not surprising because BiSkip and BiCVM require more expensive supervision at the sentence level in contrast to the other two models.

An interesting observation is that BiCCA and

BiVCD are better at separating antonyms. The words *peace* and *war*, (and their French translations *paix* and *guerre*) are well separated in BiCCA and BiVCD. However, in BiSkip and BiCVM these pairs are very close together. This can be attributed to the fact that BiSkip and BiCVM are trained on parallel sentences, and if two antonyms are present in the same sentence in English, they will also be present together in its French translation. However, BiCCA uses bilingual dictionary and BiVCD use comparable sentence context, which helps in pulling apart the synonyms and antonyms.

6 Discussion

The goal of this paper was to formulate the task of learning cross-lingual word vector representations in a unified framework, and conduct experiments to compare the performance of existing

models in a unbiased manner. We chose existing cross-lingual word vector models that can be trained on two languages at a given time. In recent work, Ammar et al. (2016) train multilingual word vectors using more than two languages; our comparison does not cover this setting. It is also worth noting that we compare here different cross-lingual word embeddings, which are not to be confused with a collection of monolingual word embeddings trained for different languages individually (Al-Rfou et al., 2013).

The paper does not cover *all* approaches that generate cross-lingual word embeddings. Some methods do not have publicly available code (Coulmance et al., 2015; Zou et al., 2013); for others, like BilBOWA (Gouws et al., 2015), we identified problems in the available code, which caused it to consistently produced results that are inferior even to mono-lingually trained vectors.¹¹ However, the models that we included for comparison in our survey are representative of other cross-lingual models in terms of the form of cross-lingual supervision required by them. For example, BilBOWA (Gouws et al., 2015) and cross-lingual Auto-encoder (Chandar et al., 2014) are similar to BiCVM in this respect. Multi-view CCA (Rastogi et al., 2015) and deep CCA (Lu et al., 2015) can be viewed as extensions of BiCCA. Our choice of models was motivated to compare different forms of supervision, and therefore, adding these models, would not provide additional insight.

7 Conclusion

We presented the first systematic comparative evaluation of cross-lingual embedding methods on several downstream NLP tasks, both intrinsic and extrinsic. We provided a unified representation for all approaches, showing them as instances of a general algorithm. Our choice of methods spans a diverse range of approaches, in that each requires a different form of supervision.

Our experiments reveal interesting trends. When evaluating on intrinsic tasks such as monolingual word similarity, models relying on cheaper forms of supervision (such as BiVCD) perform almost on par with models requiring expensive supervision. On the other hand, for cross-lingual semantic tasks, like cross-lingual document classification and dictionary induction, the model with the most informative supervision performs best

¹¹We contacted the authors of the papers and were unable to resolve the issues in the toolkit.

overall. In contrast, for the syntactic task of dependency parsing, models that are supervised at a word alignment level perform slightly better. Overall this suggests that semantic tasks can benefit more from richer cross-lingual supervision, as compared to syntactic tasks.

Acknowledgement

This material is based on research sponsored by DARPA under agreement number FA8750-13-2-0008 and Contract HR0011-15-2-0025. Approved for Public Release, Distribution Unlimited. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

- Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Polyglot: Distributed word representations for multilingual nlp. In *Proc. of CoNLL*.
- Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A Smith. 2016. Massively multilingual word embeddings. *arXiv preprint arXiv:1602.01925*.
- Francis Bond and Ryan Foster. 2013. Linking and extending an open multilingual wordnet. In *Proc. of ACL*.
- Sabine Buchholz and Erwin Marsi. 2006. Conll-X shared task on multilingual dependency parsing. In *Proc. of CoNLL*.
- Sarath Chandar, Stanislas Lauly, Hugo Larochelle, Mitesh Khapra, Balaraman Ravindran, Vikas C Raykar, and Amrita Saha. 2014. An autoencoder approach to learning bilingual word representations. In *Proc. of NIPS*.
- Jocelyn Coulmance, Jean-Marc Marty, Guillaume Wenzek, and Amine Benhalloum. 2015. Transgram, fast cross-lingual word-embeddings. In *Proc. of EMNLP*.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of ibm model 2. In *Proc. of NAACL*.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proc. of EACL*.
- Lev Finkelstein, Evgeniy Gabilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. Placing search in context: the concept revisited. In *Proc. of WWW*.
- Yoav Freund and Robert E Schapire. 1999. Large margin classification using the perceptron algorithm. *Machine learning*, 37(3):277–296.

- Stephan Gouws and Anders Søgaard. 2015. Simple task-specific bilingual word embeddings. In *Proc. of NAACL*.
- Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. Bilbowa: Fast bilingual distributed representations without word alignments. In *Proc. of ICML*.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2015. Cross-lingual dependency parsing based on distributed representations. In *Proc. of ACL*.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2016. A representation learning framework for multi-source transfer parsing. In *Proc. of AAAI*.
- David R Hardoon, Sandor Szedmak, and John Shawe-Taylor. 2004. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12):2639–2664.
- Karl Moritz Hermann and Phil Blunsom. 2014. Multilingual Models for Compositional Distributional Semantics. In *Proc. of ACL*.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2014. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *arXiv preprint arXiv:1408.3456*.
- Harold Hotelling. 1936. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. 2012. Inducing crosslingual distributed representations of words. In *Proc. of COLING*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *Proc. of MT Summit*.
- Ang Lu, Weiran Wang, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. Deep multilingual correlation for improved word embeddings. In *Proc. of NAACL*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Bilingual word representations with monolingual quality in mind. In *Proc. of the Workshop on Vector Space Modeling for NLP*.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. Universal dependency annotation for multilingual parsing. In *Proc. of ACL*.
- Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157.
- 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, Quoc V Le, and Ilya Sutskever. 2013b. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Jerome L. Myers and Arnold D. Well. 1995. *Research Design & Statistical Analysis*. Routledge.
- Pushpendre Rastogi, Benjamin Van Durme, and Raman Arora. 2015. Multiview LSA: Representation learning via generalized CCA. In *Proceedings of NAACL*.
- Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information Processing and Management*.
- Tobias Schnabel, Igor Labutov, David Mimno, and Thorsten Joachims. 2015. Evaluation methods for unsupervised word embeddings. In *Proc. of EMNLP*.
- Anders Søgaard, Željko Agić, Héctor Martínez Alonso, Barbara Plank, Bernd Bohnet, and Anders Johannsen. 2015. Inverted indexing for cross-lingual nlp. In *Proc. of ACL*.
- James H Steiger. 1980. Tests for comparing elements of a correlation matrix. *Psychological bulletin*, 87(2):245.
- O. Täckström, R. McDonald, and J. Uszkoreit. 2012. Cross-lingual word clusters for direct transfer of linguistic structure. In *NAACL*.
- Huihsin Tseng, Pichuan Chang, Galen Andrew, Daniel Jurafsky, and Christopher Manning. 2005. A conditional random field word segmenter for sghan bake-off 2005. In *Proc. of SIGHAN*.
- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Guillaume Lample, and Chris Dyer. 2015. Evaluation of word vector representations by subspace alignment. In *Proc. of EMNLP*.
- Ivan Vulić and Marie-Francine Moens. 2013a. Cross-lingual semantic similarity of words as the similarity of their semantic word responses. In *Proc. of NAACL*.
- Ivan Vulić and Marie-Francine Moens. 2013b. A study on bootstrapping bilingual vector spaces from non-parallel data (and nothing else). In *Proc. of EMNLP*.
- Ivan Vulić and Marie-Francine Moens. 2015. Bilingual word embeddings from non-parallel document-aligned data applied to bilingual lexicon induction. In *Proc. of ACL*.
- Will Y. Zou, Richard Socher, Daniel Cer, and Christopher D. Manning. 2013. Bilingual word embeddings for phrase-based machine translation. In *Proc. of EMNLP*.