

A Stronger Baseline for Multilingual Word Embeddings

Philipp Dufter, Hinrich Schütze

Center for Information and Language Processing (CIS) LMU Munich, Germany

philipp@cis.lmu.de

Abstract

Levy, Søgaard and Goldberg’s (2017) S-ID (sentence ID) method applies word2vec on tuples containing a sentence ID and a word from the sentence. It has been shown to be a strong baseline for learning multilingual embeddings. Inspired by recent work on concept based embedding learning we propose SC-ID, an extension to S-ID: given a sentence aligned corpus, we use sampling to extract concepts that are then processed in the same manner as S-IDs. We perform experiments on the Parallel Bible Corpus across 1000+ languages and show that SC-ID yields up to 6% performance increase in a word translation task. In addition, we provide evidence that SC-ID is easily and widely applicable by reporting competitive results across 8 tasks on a EuroParl based corpus.

1 Introduction

Multilingual embeddings are useful because they provide meaning representations of source and target in the same space in machine translation and because they are a basis for transfer learning.

In contrast to prior multilingual work (Zeman and Resnik, 2008; McDonald et al., 2011; Tsvetkov et al., 2014), automatically learned embeddings potentially perform as well but are more efficient and easier to use (Klementiev et al., 2012; Hermann and Blunsom, 2014b; Guo et al., 2016). Thus, multilingual word embedding learning is important for NLP.

The quality of multilingual embeddings is driven by the underlying feature set more than the type of algorithm used for training the embeddings (Upadhyay et al., 2016; Ruder et al., 2017).

Most embedding learners build on using *context information* as feature. Dufter et al. (2018) recently showed that using *concept information* can be effective for multilingual embedding learning as well.

Here, we propose the method SC-ID. SC-ID combines the concept identification method “any-malign” (Lardilleux and Lepage, 2009) and S-ID (Levy et al., 2017) into an embedding learning method that is based on both concept (C-ID) and context (S-ID). We show below that SC-ID is effective. On a massively parallel corpus, the Parallel Bible Corpus (PBC) covering 1000+ languages, SC-ID outperforms S-ID in a word translation task. In addition, we show that SC-ID outperforms S-ID on EuroParl, a corpus with characteristics quite different from PBC.

For both corpora there are embedding learners that perform better. However, SC-ID is the only one that scales easily across 1000+ languages and at the same time exhibits a stable performance across datasets. In summary, we make the following contributions in this paper: **i)** We demonstrate that using concept IDs equivalently to (Levy et al., 2017)’s sentence IDs works well. **ii)** We show that combining C-IDs and S-IDs yields higher quality embeddings than either by itself. **iii)** In extensive experiments investigating hyperparameters we find that, despite the large number of languages, lower dimensional spaces work better. **iv)** We demonstrate that our method works on very different datasets and yields competitive performance on a EuroParl based corpus.

2 Methods

Throughout this section we describe how we identify concepts and write out text corpora that are then used as input to the embedding learning algorithm.

2.1 Concept Induction

Lardilleux and Lepage (2009) propose a word alignment algorithm which we use for concept induction. They argue that hapax legomena (i.e., words which occur exactly once in a corpus) are easy to align in a sentence-aligned corpus. If hapax legomena across multiple languages occur in the same sentence, their meanings are likely the same. Similarly, words across multiple languages that occur more than once, but strictly in the same sentences, can be considered translations. We call words that occur strictly in the same sentences *perfectly aligned*. Further we define a *concept* as a set of words that are perfectly aligned.

By this definition one expects the number of identified concepts to be low. Coverage can be increased by not only considering the original parallel corpus, but by sampling subcorpora from the parallel corpus. As the number of sentences is smaller in each sample and there is a high number of sampled subcorpora, the number of perfect alignments is much higher. In addition to words, word ngrams are also considered to increase coverage. The complement of an ngram (i.e., the sentence in one language without the perfectly aligned words) is also treated as a perfect alignment as their meaning can be assumed to be equivalent as well.

For example, if for a particular subsample, the English trigram “mount of olives” occurs exactly in the same sentences as “montagne des oliviers” this gives rise to a concept, even if “olives” or “mountain” might not perfectly aligned in this particular subsample.

Figure 1 shows Lardilleux and Lepage (2009)’s anymalign algorithm. Given a sentence aligned parallel corpus, “alingual” sentences are created by concatenating each sentence across all languages. We then consider the set of all alingual sentences V , which Lardilleux and Lepage (2009) call an *alingual corpus*. The core of the algorithm iterates the following loop: (i) draw a random sample of alingual sentences $V' \subset V$; (ii) extract perfect alignments in V' . The perfect alignments are then added to the set of concepts.

Anymalign’s hyperparameters include the minimum number of languages a perfect alignment should cover (MINL) and the maximum ngram length (MAXN). The size of a subsample is adjusted automatically to maximize the probability that each sentence is sampled at least once. Ob-

Algorithm 1 Anymalign

```
1: procedure GETCONCEPTS( $V$ , MINL, MAXN, T)
2:    $C = \emptyset$ 
3:   while  $runtime \leq T$  do
4:      $V' = \text{get-subsample}(V)$ 
5:      $A = \text{get-concepts}(V')$ 
6:      $A = \text{filter-concepts}(A, \text{MINL}, \text{MAXN})$ 
7:      $C = C \cup A$ 
8:   end while
9: end procedure
```

Figure 1: V is an alingual corpus. *get-subsample* creates a subcorpus by randomly selecting lines from the alingual corpus. *get-concepts* extracts words and word ngrams that are perfectly aligned. *filter-concepts* imposes the constraint that ngrams have a specified length and concepts cover enough languages.

viously this probability depends on the number of samples drawn and thus on the runtime of the algorithm. Thus the runtime (T) is another hyperparameter. Note that T only affects the number of distinct concepts, not the quality of an individual concept. For details see (Lardilleux and Lepage, 2009).

Note that most members of a concept are neither hapax legomena nor perfect alignments in V . A word can be part of multiple concepts within V' and obviously also within V (it can be found in multiple iterations).

One can interpret the concept identification as a form of data augmentation. From an existing parallel corpus new parallel “sentences” are extracted by considering perfectly aligned subsequences.

2.2 Corpus Creation

Method S-ID. We adopt Levy and Goldberg (2014)’s framework; it formalizes the basic information that is passed to the embedding learner as a set of pairs. In the monolingual case each pair consists of two words that occur in the same context. A successful approach to multilingual embedding learning for parallel corpora is to use a corpus of pairs (one per line) of a word and a sentence ID (Levy et al., 2017). We refer to this method as *S-ID*. Note that we use S-ID for the sentence identifier, for the corpus creation method that is based on these identifiers, for the embedding learning method based on such corpora and for the embeddings produced by the method. The same applies

to C-ID. Which sense is meant should be clear from context.

Method C-ID. We use the same method for writing a corpus using our identified concepts and call this method *C-ID*. Figure 2 gives examples of the generated corpora that are passed to the embedding learner.

Method SC-ID. We combine S-IDs and C-IDs by simply concatenating their corpora before learning embeddings. However, we apply different frequency thresholds when learning embeddings from this corpus: for S-ID we investigate a frequency threshold on a development set whereas for C-ID we always set it to 1. As argued before each word in a concept carries a strong multilingual signal, which is why we do not apply any frequency filtering here. In the implementation we simply delete words in the S-ID part of the corpus with frequency lower than our threshold and set the minimum count parameter in word2vec during embedding learning to 1. Note that this is in line with how word2vec applies its minimum count threshold: words below the frequency threshold are simply ignored.

```

[...]
48001018 enge:fifteen
48001018 enge:,
48001018 enge:years
48001018 enge:after
48001018 enge:Jerusalem
48001018 enge:three
[...]
```

```

[...]
C:911 kqc0:Jerusalem
C:911 por5:Jerusalém
C:911 eng7:Jerusalem
C:911 haw0:Jerusalema
C:911 ilb0:Jelusalemu
C:911 fra1:Jérusalem
[...]
```

```

45016016 Salute one another with an holy kiss . The
churches of Christ salute you .
48001018 Then after three years I went up to
Jerusalem to see Peter , and abode with him
fifteen days .
[...]
```

Figure 2: Samples of S-ID (top) and C-ID (middle) corpora that are input to word2vec. Each word is prefixed by a 3 character ISO 639-3 language identifier followed by a alphanumeric character to distinguish multiple editions in the same language (e.g., enge = KJV). Bottom: KJV text sample. The text is whitespace tokenized.

2.3 Embedding learning

We use word2vec skipgram¹ (Mikolov et al., 2013a) with default hyperparameters, except for three: number of iterations (ITER), minimum frequency of a word (MINC) and embedding dimension (DIM). For details see Table 2.

3 Data

We work on PBC, the Parallel Bible Corpus (Mayer and Cysouw, 2014), a verse-aligned corpus of 1000+ translations of the New Testament. For the sake of comparability we use the same 1664 Bible editions across 1259 languages (distinct ISO 639-3 codes) and the same 6458 training verses as in (Dufter et al., 2018).² We follow their terminology and refer to “translations” as “editions”. PBC is a good model for resource-poverty; e.g., the training set of KJV contains fewer than 150,000 tokens in 6458 verses. KJV spans a vocabulary of 6162 words while all 32 English editions together cover 23772 unique words. See Table 1 for an example verse. We use the tokenization provided in the data, which is erroneous for some hard-to-tokenize languages (e.g., Khmer, Japanese), and do not apply additional preprocessing.

4 Evaluation

Dufter et al. (2018) introduce roundtrip translation as a multilingual embedding evaluation for languages for which no gold standard is available. We use this method for evaluating our methods on PBC. A query word w in language L_1 is translated to its closest (with respect to embedding similarity) neighbor v in L_2 and then backtranslated to its closest neighbor w' in L_1 . Roundtrip translation

¹We use code.google.com/archive/p/word2vec

²See <http://cistern.cis.lmu.de/comult>

$$q \Rightarrow I_e(q) \Rightarrow T_e(q)$$

```

woman ⇒ mujer ⇒ wife woman women widows
daughters daughter
marry married
⇒ esposa ⇒ marry wife woman married
marriage virgin daughters
bridegroom
```

Figure 3: Example for a roundtrip translation with the query “woman”. Intermediates and predictions are ordered by cosine similarity. Intermediate edition: Spanish Americas. Figure taken from (Dufter et al., 2018).

English King James Version (KJV)	German Elberfelder 1905	Spanish Americas
For the earth bringeth forth fruit of herself ; first the blade , then the ear , after that the full corn in the ear	Die Erde bringt von selbst Frucht hervor , zuerst Gras , dann eine Ähre , dann vollen Weizen in der Ähre	La tierra produce fruto por sí misma ; primero la hoja , luego la espiga , y después el grano maduro en la espiga

Table 1: Instances of verse 41004028 from the new testaments. For details on the verse ID, see (Mayer and Cysouw, 2014).

is successful if $w = w'$. As a measure of similarity we use cosine similarity. The above roundtrip is extended by considering k_I intermediate neighbors and, for each intermediate neighbor, k_T predictions in the query language.

The predictions are then compared to a groundtruth set G_w . There is a strict and a relaxed groundtruth. The former only contains the query word w whereas the latter contains all words with the same lemma and part-of-speech as w . This accommodates the fact that for a single query multiple translations can be considered correct (e.g., an inflected form of a query). We average the binary results (per query and intermediate edition) over editions and report the mean and median over queries. Inspired by the precision@k evaluation measure for word translation we vary (k_I, k_T) as follows: “S1” (1, 1), “R1” (1, 1), “S4” (2, 2), and “S16” (2, 8), where S stands for using the strict groundtruth, R for using the relaxed groundtruth.

If w is not contained in the embedding space, then the predictions are counted as incorrect. The number of queries contained in the embedding space is denoted by N . We use the same queries as in (Dufter et al., 2018), which are based on (Swadesh, 1946)’s 100 universal words. As we perform hyperparameter selection we introduce a development set of queries: an earlier list by Swadesh that contains 215 words.³ After excluding the queries from the test set, there are 151 queries in the development set. Due to this large number of queries we do not compute the relaxed measure as this requires manual creation of the groundtruth. We work on the level of Bible editions, i.e., two editions in the same language are considered different “languages”. We use KJV as the query edition if KJV contains the query word. Else we randomly choose another English edition.

5 Baselines

We compute a diverse set of baselines to pinpoint reasons for performance changes as much as possible.

³See concepticon.clld.org/contributions/Swadesh-1950-215

5.1 Context Based Embedding Space

We use (Levy et al., 2017)’s method and call it **S-ID**.

5.2 Transformation Based Embedding Space

The baseline **LINEAR** follows (Duong et al., 2017). We pick one edition as the “embedding space defining edition”, in our case we pick English KJV. We then create 1664 bilingual embedding spaces using S-ID; in each case the two editions covered are KJV and one of the other 1664 languages. We then use (Mikolov et al., 2013b)’s linear transformation method to map all embedding spaces to the KJV embedding space. More specifically, let $X \in \mathbb{R}^{n_X \times d}$, $Y \in \mathbb{R}^{n_Y \times d}$ be two embedding spaces, with n_X , n_Y number of words, and d be the embedding dimension. We then select n_T transformation words that are contained in both embedding spaces. This gives us $X' \in \mathbb{R}^{n_T \times d}$, $Y' \in \mathbb{R}^{n_T \times d}$. The transformation matrix $W \in \mathbb{R}^{d \times d}$ is then given by

$$\arg \min_W \|X'W - Y'\|_F$$

where $\|\cdot\|_F$ denotes the Frobenius norm. The closed form solution for the transformation is given by

$$W^* = X^+Y$$

where X^+ is the Moore-Penrose Pseudo inverse (Penrose, 1956). In our case X and Y are bilingual embedding spaces where both contain the vocabulary of the English KJV edition.

5.3 Bilingual Embedding Spaces

In **BILING** we use the same bilingual embedding spaces as in **LINEAR**. However, we perform roundtrip translation in each embedding space separately, i.e., there is no transformation to a common space. This baseline allows us to assess the effect of a massively multilingual compared to having only bilingual embedding spaces.

5.4 Unsupervised Embedding Learning

We apply the recent unsupervised embedding learning method by Lample et al. (2018) and call

it **MUSE**. Given unaligned corpora in two languages, MUSE learns two separate embedding spaces that are subsequently unified by a linear transformation. This transformation is learned using a discriminator neural network that tries to identify the original language of a word vector. Subsequently a refinement step based on the Procrustes algorithm is performed. Note that this method only yields bilingual embedding spaces. As in **BILING**, we perform roundtrip translation in each embedding space separately.

Very recently, [Chen and Cardie \(2018\)](#) extended MUSE multilingually. We will include this in future work.

5.5 Non-Embedding-Space-Based Baseline

To show that the embedding space provides some advantages over just using the concepts as is, we introduce **C-SIMPLE**, a non-embedding baseline that follows the idea of roundtrip translation. Given a query word and an intermediate edition we consider all words that share a concept ID with the query word as possible intermediate words. We then choose randomly (probability weights according to their relative frequency across concept) intermediate words. For the back translation we apply the same procedure. This baseline is inspired by ([Dufter et al., 2018](#))’s **RTSIMPLE** baseline.

6 Results

6.1 Hyperparameters

Word2vec. Since a grid search for optimal values for the parameters ITER, MINC and DIM would take too long, we search greedily instead. More iterations yield better performance. ITER = 100 is a good efficiency-performance trade-off. For MINC (minimum frequency of a word) the best performance is found using 2. This is mainly due to increased coverage. Surprisingly, smaller embedding dimensions work better to some degree. A highly multilingual embedding space is expected to suffer more from ambiguity and that is an argument for higher dimensionality; cf. [Li and Jurafsky \(2015\)](#). But this effect seems to be counteracted by the low-resource properties of PBC for which the increased number of parameters of higher dimensionalities cannot be estimated reliably. We choose embedding size 100.

C-ID. We use the tuned word2vec settings to investigate the hyperparameters for C-ID. As men-

	ITER	MINC	DIM	S1		S4		S16		N	
				μ	Md	μ	Md	μ	Md		
1	S-ID	<i>100</i>	5	200	29	21	43	46	56	78	103
2	S-ID	5			14	11	25	22	41	45	103
3	S-ID	10			25	16	38	34	51	60	103
4	S-ID	25			27	20	41	43	53	69	103
5	S-ID	50			27	16	40	40	53	67	103
6	S-ID	150			29	21	43	47	56	79	103
7	S-ID		2		35	31	52	60	66	90	130
8	S-ID		10		24	5	36	17	48	63	85
9	S-ID			<i>100</i>	30	24	45	48	58	83	103
10	S-ID			300	28	19	42	41	54	69	103

	MINL	MAXN	T	S1		S4		S16		N	
				μ	Md	μ	Md	μ	Md		
1	C-ID	<i>100</i>	3	10	30	26	46	46	60	71	120
2	C-ID	50			26	21	39	42	56	79	104
3	C-ID	150			22	14	34	28	47	56	94
4	C-ID	500			15	0	25	0	33	0	59
5	C-ID		1		17	0	27	0	38	0	73
6	C-ID			5	28	20	40	35	55	58	122
7	C-ID			5	24	20	36	34	48	52	114
8	C-ID			15	30	26	46	44	61	76	121

Table 2: Hyperparameter optimization for word2vec (top) and anymalign (bottom). Initial parameters in first row; empty cell = unchanged parameter. Bold: best result per column. Selected hyperparameter values: *italics*.

tioned we set MINC to 1 whenever we learn on a concept based corpus. MINL is the minimum number of editions. We find the best performance when setting MINL to 100. As expected coverage worsens when choosing higher values for MINL. MAXN is the maximum ngram length. We see the best performance when setting MAXN to 3. This seems intuitive, as this allows for – at least in European languages – common word compounds. T is the time in hours, i.e., for how long to sample (which inherently steers the size and number of sample subcorpora). For T, 10 and 15 differ only slightly. We set T to 15 to get slightly higher coverage. See Table 3 for basic statistics about the identified concepts.

6.2 Roundtrip Translation Results

Table 4 presents roundtrip translation results on the test set.

C-Simple performs reasonably well, but is outperformed by almost all embedding spaces. This indicates that learning embeddings augments the information stored in concepts. LINEAR per-

	Mean	Median	Std.	Max	Min
#editions	250	194	160	101	1530
#tokens	259	198	172	102	2163

Table 3: Descriptive statistics of concept size for our selected hyperparameters. On average a concept contains 259 tokens across 250 editions. The total number of concepts is 119,040. The largest concept contains 2163 tokens across 1530 editions describing the 2-gram ‘‘Simon Peter’’.

forms similar to C-Simple (but outperforms it for S16). This supports the hypothesis that PBC offers too little data for training mono-/bilingual embeddings. As expected BILING works better than LINEAR. However, keep in mind that this is not a universal embedding space and has fewer constraints than the other embedding spaces. Thus it is not directly comparable. MUSE works surprisingly well given that this method does not exploit the sentence alignment information in the corpus. While this method is ranked last in S1 and R1 (μ) it gains traction on S4 and S16 where it ranks sixth in S16. However, the numbers are not directly comparable as no multilingual embedding space is created by MUSE (only 1664 bilingual ones). Thus there are fewer constraints overall. The performance drop between BILING and LINEAR can be considered as rough proxy for how much we would expect MUSE to worsen if considering a true multilingual embedding space. In future work we want to investigate this effect using (Chen and Cardie, 2018). Note that MUSE is quite inefficient: it takes around 1 hour to train a bilingual embedding space for a language pair on a standard GPU.

N(t) by Dufter et al. (2018) shows consistently the best performance. However, we will see later that this method only works in a massively multilingual setting and fails on EuroParl. SAMPLE by Dufter et al. (2018) is based on sampling like C-ID. However, it induces concepts only on a small set of pivot languages, not on all 1664 editions. This does not only work worse (C-ID beats SAMPLE except for S16), but it also requires additional word alignment information and is thus computationally more expensive.

As has been observed frequently, S-ID is highly effective. S-ID ranks consistently third. Representing a word as its binary vector of verse occurrence provides a clear signal to the embedding learner. Concept-based methods can be consid-

	S1		R1		S4		S16		N
	μ	Md	μ	Md	μ	Md	μ	Md	
1 C-SIMPLE	35	33	35	34	49	54	56	56	67
2 S-ID	48	47	53	59	65	72	83	93	69
3 C-ID	43	42	46	43	58	60	79	91	67
4 SC-ID	51	50	56	65	69	80	86	96	69
5 LINEAR	35	32	36	34	47	51	65	72	69
6 BILING [†]	42	34	44	41	55	54	70	85	69
7 MUSE [†]	31	35	31	35	52	57	78	86	69
8 N(t)*	54	59	61	69	80	87	94	100	69
9 SAMPLE*	33	23	43	42	54	59	82	96	65

Table 4: Roundtrip evaluation of multilingual embedding learning methods. Results marked with * are taken from (Dufter et al., 2018). Results marked with [†] are bilingual embedding spaces.

ered complementary to this feature set because they can exploit aggregated information across languages as well as across verses. C-ID alone has slightly lower performance. SC-ID, the combination, outperforms C-ID and S-ID and seems to be the ‘‘best of both worlds’’. It yields consistently the best performance with 3% to 6% relative performance increase (for μ) compared to S-ID alone and always ranks on the second place. In short: N(t) is the best method on this corpus followed by SC-ID and subsequently S-ID.

7 Application to a High-Resource Corpus

PBC is a good model for learning multilingual embeddings in a low-resource, highly multilingual scenario. We now provide experimental evidence that SC-ID is broadly applicable and also works in a high-resource, mildly multilingual scenario. We test the three best performing methods (based on PBC S1 μ): S-ID, SC-ID and N(t).

7.1 Data

We choose a dataset published by Ammar et al. (2016),⁴ a parallel corpus covering 12 languages (Bulgarian, Czech, Danish, English, Finnish, French, German, Greek, Hungarian, Italian, Spanish, Swedish) from the proceedings of the European Parliament, Wikipedia titles and parallel news commentary. We refer to this corpus as *EuroParl*.

7.2 Evaluation

Ammar et al. (2016) provide an extensive evaluation framework covering one extrinsic (their dependency parsing models are not available for

⁴See <http://128.2.220.95/multilingual/data/>

	Document Classification		Word Similarity		QVEC		QVEC-CCA		Word Translation							
	monol.	multil.	monol.	multil.	monol.	multil.	monol.	multil.	monol.	multil.						
multiCluster*	92.11	48.16	38.07	57.57	57.45	73.89	10.36	98.62	9.32	82.01	62.46	98.62	43.34	82.01	43.94	45.22
multiCCA*	92.18	62.81	43.09	71.09	69.99	77.94	10.75	99.09	8.78	87.03	63.50	99.09	41.52	87.03	36.21	54.87
multiSkip*	90.46	45.73	33.94	55.41	60.24	67.55	8.41	98.09	7.21	75.69	58.98	98.09	36.34	75.69	46.71	39.55
INVARIANCE*	91.10	31.35	51.08	23.06	59.13	62.50	8.11	91.71	5.38	74.78	65.83	91.71	46.21	74.78	63.91	30.36
S-ID	86.47	38.89	47.12	36.04	53.26	56.46	12.59	94.19	11.82	71.00	34.63	94.19	25.86	71.00	61.64	33.89
SC-ID	86.99	46.43	45.71	46.41	53.94	66.27	12.36	97.05	11.07	75.34	34.67	97.05	26.04	75.34	64.99	38.72

Table 5: RIGHT: Word translation results on the test dataset. Methods marked with star are from (Ammar et al., 2016).

LEFT: Results on several tasks on the test set. We downloaded their embedding spaces and performed the evaluation using their code. While almost all results are reproduced (except for rounding) we get slightly different results in the multilingual word similarity task and for multilingual QVEC (multiskip only).

SC-ID		S-ID		N(t)			
MINL	MINC	MINC	Word Trans.	NPIV	MINC	Word Trans.	
6	5	63.86	37.51	4	5	57.14	0.65
3		63.03	37.42	2		42.31	2.41
9		61.88	37.51	6		33.33	0.28
12		60.15	37.51	1		42.86	0.65
2		57.66	48.47	10		42.86	0.65
10		60.97	36.40				
				5		59.90	37.51
				2		54.46	46.89
				10		62.02	31.29

Table 6: Hyperparameter selection on EuroParl. Initial parameter in first row; empty cell = unchanged parameter. Bold: best result per column. Selected hyperparameter values: *italics*. Subscript numbers indicate the coverage.

download, so we omit this part of the evaluation) and four intrinsic tasks that we use in this work. The tasks are document classification, word translation, word similarity, QVEC (Tsvetkov et al., 2015) and QVEC-CCA (Ammar et al., 2016). In addition there are 3 monolingual tasks (word similarity, QVEC and QVEC-CCA) that are all performed in English. Our main focus remains on word translation (precision@1 is reported). For all tasks there is a development and test set available. For a detailed description of the evaluation data and tasks see (Ammar et al., 2016).

To ensure comparability with previous approaches we follow Ammar et al. (2016) in evaluating only on words that are contained in the embedding space and then simultaneously reporting the coverage (i.e., how many words of the translation task are actually contained in the embedding space). Note that this is different from our procedure on PBC. Throughout, subscript numbers in tables indicate coverage in percent.

7.3 Hyperparameters

We optimize corpus specific hyperparameters (e.g., number of languages a concept should contain) on the development set of the word translation task.

SC-ID: we vary MINC and the minimum number of languages that need to be covered by the concept identification algorithm. For S-ID we only vary MINC. N(t) requires pivot languages. Following Dufter et al. (2018) we choose as pivot languages those with the lowest type-token ratio (these are Greek, Danish, Spanish, French, Italian, English) and vary the number between 2, 4 and 6.

Table 6 gives an overview of our hyperparameter selection. For SC-ID, we choose MINL 6 and MINC 5 (compared to 2 in PBC). Given that this corpus provides much more information per language it seems intuitive that MINC should be higher (again note that for the C-ID part we always use MINC 1). This is confirmed for S-ID where we find the best performance with MINC 10. For N(t) the best result is obtained when using 4 pivot languages. N(t) exhibits an exceptionally low coverage, which obviously varies with the number of pivot languages: more pivot languages make it less likely that N(t) will find two words that have exactly the same neighbors in the dictionary graph. See (Dufter et al., 2018) for details on N(t). Even for 2 pivot languages the coverage is too low to be useful for any application. We computed the coverage on the test set, which is only 0.56 when using 4 pivot languages. Still the precision is quite good. This indicates that N(t) is an effective method, but is only applicable to a corpus covering a large number of languages. Thus we do not report results on the test set.

7.4 Results

Table 5 (RIGHT) gives results on the word translation test set. SC-ID performs best followed by INVARIANCE (Huang et al., 2015). While INVARIANCE is theoretically applicable to PBC, there are strong indications why this is not possible in practice. INVARIANCE considers the full cooccurrence matrix across all languages, a matrix that would be in the size of terabytes. In addition, word alignment matrices would need to be stored. The methods proposed by Ammar et al. (2016) yield reasonable, but clearly lower performance.

Table 5 (LEFT) gives results for the remaining tasks. Immediately it becomes clear that different embedding spaces have different strengths and weaknesses. The results for SC-ID are competitive throughout all tasks and there is no method which consistently outperforms SC-ID. SC-ID beats S-ID in 4 out of 7 tasks and in 3 out of 4 multilingual tasks. In addition SC-ID provides a significantly higher coverage throughout all tasks.

Overall SC-ID outperforms S-ID clearly on PBC and on word translation on EuroParl. On other tasks it outperforms S-ID in three out of four multilingual tasks and else has a competitive performance. None of the many baselines in this paper consistently outperforms SC-ID on both datasets. Thus, SC-ID is a simple baseline, yet a stronger one than S-ID.

8 Related Work

Much research has been dedicated to identifying **multilingual concepts**. BabelNet (Navigli and Ponzetto, 2012) leverages existing resources (mostly manual annotations), including Wikipedia using information extraction methods. While BabelNet could be directly used to learn concept based embeddings, it covers only 284 languages and thus cannot be applied to all PBC languages. Other work induces concepts within a dictionary graph (Ammar et al., 2016; Dufter et al., 2018) or with alignment algorithms (Östling, 2014).

We now review **multilingual embedding learning** methods for parallel corpora. We have clustered them into three groups.

Group 1 consists of using monolingual algorithms for creating monolingual embedding spaces. Subsequently they are projected into a unified space using a (linear) transforma-

tion. We use (Mikolov et al., 2013b) together with (Duong et al., 2017) in our baseline LINEAR. Zou et al. (2013), Xiao and Guo (2014) and Faruqui and Dyer (2014) use similar approaches (e.g., by computing the transformation using CCA). Recently it has been shown that computing the transformation using discriminator neural networks works well, even in a completely unsupervised setting. See, e.g., (Vulić and Moens, 2012; Ammar et al., 2016; Lample et al., 2018; Chen and Cardie, 2018). We used (Lample et al., 2018) as a baseline.

Group 2 is true multilingual embedding learning: it integrates multilingual information in the objective of embedding learning. Klementiev et al. (2012) and Gouws et al. (2015) add a word alignment based term. Luong et al. (2015) introduce BiSkip as a bilingual extension of word2vec. For n editions, including $O(n^2)$ bilingual terms does not scale. A slightly different objective function expresses that representation of aligned sentences should be similar. Approaches based on neural networks are (Hermann and Blunsom, 2014a) (BiCVM), (Sarath Chandar et al., 2014) (autoencoders) and (Soyer et al., 2014).

Group 3 creates multilingual corpora and uses monolingual embedding learners. A successful approach is (Levy et al., 2017)’s sentence ID (S-ID). Vulić and Moens (2015) create pseudo-corpora by merging words from multiple languages into a single corpus. Dufter et al. (2018) found this method to perform poorly on PBC. Sjøgaard et al. (2015) learn a space by factorizing an interlingual matrix based on Wikipedia concepts. word2vec is roughly equivalent to matrix factorization (Levy and Goldberg, 2014), so this work fits this category.

9 Summary

S-ID, the method for multilingual embedding learning proposed by Levy et al. (2017), is a surprisingly strong baseline. It exploits *sentential context of words*. In this paper, we showed that what we call C-ID – concepts learned by Lardilleux and Lepage (2009)’s method – is a competitive alternative. Concepts exploit the *conceptual content of words*. We demonstrated that SC-ID, the *combination of concept information and context information*, performs better by a large margin than each by itself. We provided ex-

perimental evidence for this conclusion on two datasets with different characteristics: one low-resource and highly multilingual (PBC) and one high-resource and mildly multilingual (EuroParl).

Acknowledgements. We gratefully acknowledge funding for this work by the European Research Council (ERC #740516) and by Zentrum Digitalisierung.Bayern (ZD.B), the digital technology initiative of the State of Bavaria.

References

- Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A Smith. 2016. Massively multilingual word embeddings. *arXiv preprint arXiv:1602.01925*.
- Xilun Chen and Claire Cardie. 2018. Unsupervised multilingual word embeddings. *arXiv preprint arXiv:1808.08933*.
- Philipp Dufter, Mengjie Zhao, Martin Schmitt, Alexander Fraser, and Hinrich Schütze. 2018. Embedding learning through multilingual concept induction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Long Duong, Hiroshi Kanayama, Tengfei Ma, Steven Bird, and Trevor Cohn. 2017. Multilingual training of crosslingual word embeddings. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*.
- Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. Bilbowa: fast bilingual distributed representations without word alignments. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning*.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2016. A representation learning framework for multi-source transfer parsing. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*.
- Karl Moritz Hermann and Phil Blunsom. 2014a. Multilingual distributed representations without word alignment. In *Proceedings of the 2014 International Conference on Learning Representations*.
- Karl Moritz Hermann and Phil Blunsom. 2014b. Multilingual models for compositional distributed semantics. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*.
- Kejun Huang, Matt Gardner, Evangelos Papalexakis, Christos Faloutsos, Nikos Sidiropoulos, Tom Mitchell, Partha P Talukdar, and Xiao Fu. 2015. Translation invariant word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1084–1088.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. 2012. Inducing crosslingual distributed representations of words. *Proceedings of the 24th International Conference on Computational Linguistics*.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *Proceedings of the 2018 International Conference on Learning Representations*.
- Adrien Lardilleux and Yves Lepage. 2009. Sampling-based multilingual alignment. In *Proceedings of 7th Conference on Recent Advances in Natural Language Processing*.
- Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*.
- Omer Levy, Anders Søgaard, and Yoav Goldberg. 2017. A strong baseline for learning cross-lingual word embeddings from sentence alignments. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*.
- Jiwei Li and Dan Jurafsky. 2015. Do multi-sense embeddings improve natural language understanding? In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Bilingual word representations with monolingual quality in mind. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*.
- Thomas Mayer and Michael Cysouw. 2014. Creating a massively parallel bible corpus. In *Proceedings of the 9th International Conference on Language Resources and Evaluation*.
- Ryan T. McDonald, Slav Petrov, and Keith B. Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*.
- 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*.
- Roberto Navigli and Simone Paolo Ponzetto. 2012. BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*.
- Robert Östling. 2014. Bayesian word alignment for massively parallel texts. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*.
- Roger Penrose. 1956. On best approximate solutions of linear matrix equations. In *Mathematical Proceedings of the Cambridge Philosophical Society*.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2017. A survey of cross-lingual word embedding models. *arXiv preprint arXiv:1706.04902*.

- AP Sarath Chandar, Stanislas Lauly, Hugo Larochelle, Mitesh M. Khapra, Balaraman Ravindran, Vikas C. Raykar, and Amrita Saha. 2014. An autoencoder approach to learning bilingual word representations. In *Proceedings of the 2014 Annual Conference on Neural Information Processing Systems*.
- Anders Søgaard, Željko Agić, Héctor Martínez Alonso, Barbara Plank, Bernd Bohnet, and Anders Johannsen. 2015. Inverted indexing for cross-lingual nlp. In *The 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference of the Asian Federation of Natural Language Processing*.
- Hubert Soyer, Pontus Stenetorp, and Akiko Aizawa. 2014. Leveraging monolingual data for crosslingual compositional word representations. In *Proceedings of the 2015 International Conference on Learning Representations*.
- Morris Swadesh. 1946. South Greenlandic (Eskimo). In Cornelius Osgood, editor, *Linguistic Structures of Native America*, Viking Fund Inc. (Johnson Reprint Corp.), New York.
- Yulia Tsvetkov, Leonid Boytsov, Anatole Gershman, Eric Nyberg, and Chris Dyer. 2014. Metaphor detection with cross-lingual model transfer. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*.
- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Guillaume Lample, and Chris Dyer. 2015. Evaluation of word vector representations by subspace alignment. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2049–2054.
- Shyam Upadhyay, Manaal Faruqui, Chris Dyer, and Dan Roth. 2016. Cross-lingual models of word embeddings: An empirical comparison. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Ivan Vulić and Marie-Francine Moens. 2012. Detecting highly confident word translations from comparable corpora without any prior knowledge. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*.
- Ivan Vulić and Marie-Francine Moens. 2015. Bilingual word embeddings from non-parallel document-aligned data applied to bilingual lexicon induction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*.
- Min Xiao and Yuhong Guo. 2014. Distributed word representation learning for cross-lingual dependency parsing. In *Proceedings of the 18th Conference on Computational Natural Language Learning*.
- Daniel Zeman and Philip Resnik. 2008. Cross-language parser adaptation between related languages. In *Proceedings of the 3rd International Joint Conference on Natural Language Processing*.
- Will Y Zou, Richard Socher, Daniel Cer, and Christopher D Manning. 2013. Bilingual word embeddings for phrase-based machine translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*.