

Language classification from bilingual word embedding graphs

Steffen Eger
UKP Lab
TU Darmstadt
64289 Darmstadt

Armin Hoenen
Text Technology Lab
Goethe University
60325 Frankfurt am Main

Alexander Mehler
Text Technology Lab
Goethe University
60325 Frankfurt am Main

Abstract

We study the role of the second language in bilingual word embeddings in monolingual semantic evaluation tasks. We find strongly and weakly positive correlations between down-stream task performance and second language similarity to the target language. Additionally, we show how bilingual word embeddings can be employed for the task of semantic language classification and that joint semantic spaces vary in meaningful ways across second languages. Our results support the hypothesis that semantic language similarity is influenced by both structural similarity as well as geography/contact.

1 Introduction

Word embeddings derived from context-predicting neural network architectures have become the state-of-the-art in distributional semantics modeling (Baroni et al., 2014). Given the success of these models and the ensuing hype, several extensions over the standard paradigm (Bengio et al., 2003; Collobert and Weston, 2008; Mikolov et al., 2013; Pennington et al., 2014) have been suggested, such as retrofitting word vectors to semantic knowledge-bases (Faruqui et al., 2015), multi-sense (Huang et al., 2012; Neelakantan et al., 2014), and multi-lingual word vectors (Klementiev et al., 2012; Faruqui and Dyer, 2014; Hermann and Blunsom, 2014; Chandar et al., 2014; Lu et al., 2015; Gouws et al., 2015; Gouws and Søgaard, 2015; Huang et al., 2015; Šuster et al., 2016).

The models underlying the latter paradigm, which we focus on in the current work, project word vectors of two (or multiple) languages into a joint semantic space, thereby allowing to evaluate semantic similarity of words from different languages; see Figure 1 for an illustration. Moreover, the resulting word vectors have been shown to produce on-par or better performance even in a *monolingual* setting, e.g., when using them for measuring semantic similarity in one of the two languages involved (Faruqui and Dyer, 2014).

While multilingual word vectors have been evaluated with respect to intrinsic parameters such as embedding dimensionality, empirical work on another aspect appears to be lacking: the second language involved. For example, it might be the case that projecting two languages with very different lexical semantic associations in a joint embedding space inherently deteriorates *monolingual* embeddings as measured by performance on an intrinsic monolingual semantic evaluation task, relative to a setting in which the two languages have very similar lexical semantic associations. To illustrate, the classical Latin word *vir* is sometimes translated in English as both ‘man’ and ‘warrior’, suggesting a semantic connotation, in Latin, that is putatively lacking in English. Hence, projecting English and Latin in a joint semantic space may invoke semantic relations that are misleading for an English evaluation task. Alternatively, it may be argued that heterogeneity in semantics between the two languages involved is beneficial for monolingual evaluation tasks in the same way that uncorrelatedness in classifiers helps in combining them.

Here, we study two questions (**main contributions**). On the one hand, we are interested in the effect of language similarity on bilingual word embeddings in a (Q1) **monolingual (intrinsic) semantic evaluation task**. Thus, our first question is: how does the performance of bilingual word embeddings

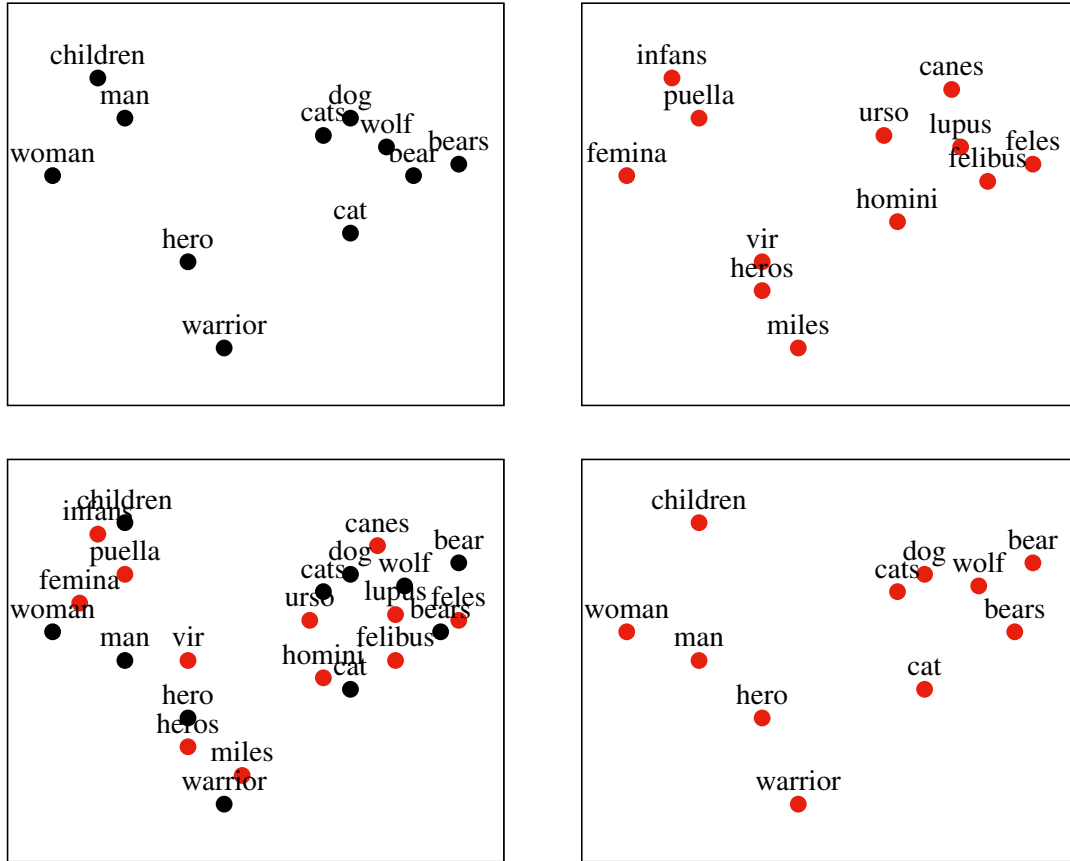


Figure 1: Monolingual embeddings (top left and top right) have been shown to capture semantic (as well as syntactic) properties of languages, here exemplarily: p = English and ℓ = Latin. Bottom left: The (idealized) goal of crosslingual embeddings is to capture these relationships across two or more languages. Bottom right: After projection in a joint semantic space, semantic (as well as syntactic) properties of words in language p have adapted to those of language ℓ . Note, in particular, the movement of *man* in these idealized plots, i.e., the different positions of *man* in top left vs. bottom right.

in monolingual semantic evaluation tasks depend on the second language involved?¹ Secondly, we ask how bilingual word embeddings can be employed for the task of semantic (Q2) **language classification**. Our approach here is simple: we project languages onto a common pivot language p so as to make them comparable. We directly use bilingual word embeddings for this. More precisely, we first project languages ℓ in a common semantic space with the pivot p by means of bilingual word embedding methods. Subsequently, we ignore language ℓ words in the joint space. Semantic distance measurement between languages then amounts to comparison of graphs that have the same nodes — pivot language words — and different edge weights — semantic similarity scores between pivot language words based on bilingual embeddings that vary as a function of the second language ℓ involved. This core idea is illustrated in Figures 1 and 2.

We show that joint semantic spaces induced by bilingual word embeddings vary in meaningful ways across second languages. Moreover, our results support the hypothesis that semantic language similarity is influenced by both genealogical language similarity and by aspects of language contact.

This work is structured as follows. Section 2 introduces our approach of constructing graphs from bilingual word embeddings and its relation to the two questions outlined. Section 3 describes our data,

¹Our initial expectation was that bilingual word embeddings lead to better results in monolingual settings, at least for some second languages. However, this was not confirmed in any of our experiments. This may be related to our (small) data set sizes (see Section 3) or to other factors, but has no concern for the question (Q1) we are investigating.

which is based on the Europarl corpus (Koehn, 2005). Section 4 details our experiments, which we discuss in Section 5. We relate to previous work in Section 6 and conclude in Section 7.

2 Model

In this section, we formally outline our approach.

Given $N + 1$ languages, choose one of them, p , as *pivot language*. Construct N weighted networks $G_\ell^{(p)} = (V^{(p)}, E^{(p)}, w_\ell^{(p)})$ as follows: nodes $V^{(p)}$ are the words of language p , graphs are fully connected, i.e., $E^{(p)} = V^{(p)} \times V^{(p)}$, and edge weights are $w_\ell^{(p)}(u, v) = \text{sim}(\mathbf{u}_{p,\ell}, \mathbf{v}_{p,\ell})$. The similarity function sim is, e.g., cosine similarity, and $\mathbf{u}_{p,\ell}, \mathbf{v}_{p,\ell} \in \mathbb{R}^d$ are *bilingual* word embeddings of words u and v , respectively, derived from any suitable method (see below). Here, ℓ ranges over the N *second languages*.

For (Q1) **monolingual semantic evaluation** in language p , choose p as pivot and consider $G_\ell^{(p)}$ for varying second languages ℓ . We can then evaluate semantic similarity between two language p words u and v by querying the edge weight $w_\ell^{(p)}(u, v)$. This is the classical situation of (intrinsic) monolingual evaluation of bilingual word embeddings.

For (Q2) **language classification**, we compare the graphs $G_\ell^{(p)}$ across all second languages ℓ , and a fixed pivot p . Here, we have many choices how to realize distance measures between graphs, such as which metric we use and at which level we compare graphs (Bunke and Shearer, 1998; Rothe and Schütze, 2014). We choose the following: we first represent each node (pivot language word) in a graph as the vector of distances to all other nodes. That is, to $u \in V[G_\ell^{(p)}] = V^{(p)}$ we assign the vector $\mathbf{r}_u^{(\ell)} = (w_\ell^{(p)}(u, v))_{v \in V^{(p)}}$. The distance $d(G_\ell^{(p)}, G_{\ell'}^{(p)})$ between two graphs $G_\ell^{(p)}$ and $G_{\ell'}^{(p)}$ is then defined as the average distance (Euclidean norm) of the so represented nodes in the graphs: $d(G_\ell^{(p)}, G_{\ell'}^{(p)}) = \frac{1}{|V^{(p)}|} \sum_{u \in V^{(p)}} \|\mathbf{r}_u^{(\ell)} - \mathbf{r}_u^{(\ell')}\|$. Finally, we define the (*syntacto-*)*semantic distance* $D(\ell, \ell')$ of two languages ℓ and ℓ' as the average graph distance over all $N - 1$ pivots (ℓ and ℓ' excluded):

$$D(\ell, \ell') = \frac{1}{N-1} \sum_{\tilde{p}} d(G_\ell^{(\tilde{p})}, G_{\ell'}^{(\tilde{p})}). \quad (1)$$

By summing over pivots, we effectively ‘integrate out’ the influence of the pivot language, leading to a ‘pivot independent’ language distance calculation. In addition, this ensures that the distance matrix D encompasses all languages, including all possible pivots.

Figure 2 illustrates our idea of projecting semantic spaces of different languages onto a common pivot.




Figure 2: Schematic illustration of our approach. Repeating the “four-stage” process illustrated in Figure 1 for three different languages ℓ (marked by different colors) and the same pivot p . Edge strengths between pivot language words indicate their semantic similarity as measured by cosine distances in semantic spaces as in Figure 1 bottom right.

Bilingual embedding models: We consider two approaches to constructing bilingual word embeddings. The first is the canonical correlation analysis (CCA) approach suggested in Faruqi and Dyer

(2014). This takes *independently* constructed word vectors from two different languages and projects them onto a common vector space such that (one-best) translation pairs, as determined by automatic word alignments, are maximally linearly correlated. CCA relies on word level alignments and we use cdec for this (Dyer et al., 2010).

The second approach we employ is called BilBOWA (**BBA**) (Gouws et al., 2015). Rather than separately training word vectors for two languages and subsequently enforcing cross-lingual constraints, this model *jointly* optimizes monolingual and cross-lingual objectives similarly as in Klementiev et al. (2012):

$$\mathcal{L} = \sum_{\ell \in \{e, f\}} \sum_{w, h \in \mathcal{D}^\ell} \mathcal{L}^\ell(w, h; \theta^\ell) + \lambda \Omega(\theta^e, \theta^f)$$

is minimized, where w and h are target words and their contexts, respectively, and θ^e, θ^f are embedding parameters for two languages. The terms \mathcal{L}^ℓ encode the monolingual constraints and the term $\Omega(\theta^e, \theta^f)$ encodes the cross-lingual constraints, enforcing similar words across languages (obtained from *sentence aligned* data) to have similar embeddings.

3 Data

For our experiments, we use the Wikipedia extracts available from Al-Rfou et al. (2013)² as monolingual data and Europarl (Koehn, 2005) as bilingual database. We consider two settings, one in which we take all 21 (All21) languages available in Europarl and one in which we focus on the 10 (Big10) largest languages. These languages are bg, cs, **da, de**, el, **en, es**, et, **fi, fr**, hu, **it**, lt, lv, **nl, pl, pt**, ro, sk, sl, **sv** (Big10 languages highlighted). To induce a comparable setting, we extract in the All21 setup: 195,842 parallel sentences from Europarl and roughly 835K (randomly extracted) sentences from Wikipedia for each of the 21 languages. In the Big10 setup, we extract 1,098,897 parallel sentences from Europarl and 2,540K sentences from Wikipedia for each of the 10 languages involved. We note that the above numbers are determined by the minimum available for the respective two sets of languages in the Europarl and Wikipedia data, respectively. As preprocessing, we tokenize all sentences in all datasets and we lower-case all words.

4 Experiments

We first train $d = 200$ dimensional skip-gram word2vec vectors (Mikolov et al., 2013)³ on the union of the Europarl and Wikipedia data for each language in the respective All21 and Big10 setting. For CCA, we then obtain bilingual embeddings for each possible combination (ℓ, ℓ') of languages in each of the two setups, by projecting these vectors in a joint space via word alignments obtained on the respective Europarl data pair. For BBA, we use the monolingual Wikipedias of ℓ and ℓ' for the monolingual constraints, and the Europarl sentence alignments of ℓ and ℓ' for the bilingual constraints. We only consider words that occur at least 100 times in the respective data sets.

4.1 Monolingual semantic task (Q1)

We first evaluate the obtained BBA and CCA embedding vectors on monolingual $p = \text{English}$ evaluation tasks, for varying second language ℓ . The tasks we consider are WS353 (Finkelstein et al., 2002), MTurk287 (Radinsky et al., 2011), MTurk771,⁴ SimLex999 (Hill et al., 2015), and MEN (Bruni et al., 2014), which are standard semantic similarity datasets for English, documented in an array of previous research. In addition, we include the SimLex999-de and SimLex999-it (Leviant and Reichart, 2015) for $p = \text{German}$ and $p = \text{Italian}$, respectively. In each task, the goal is to determine the semantic similarity between two language p words, such as *dog* and *cat* (when $p = \text{English}$). For the tasks, we indicate average Spearman correlation coefficients $\delta = \delta_{p, \ell}$ between

²<https://sites.google.com/site/rmyeid/projects/polyglot>.

³All other parameters set to default values.

⁴<http://www2.mta.ac.il/gideon/mturk771.html>

- the predicted semantic similarity — measured in cosine similarity — between respective language p word pair vectors obtained when projecting p and ℓ in a joint embedding space, and
- the human gold standard (i.e., human judges have assigned semantic similarity scores to word pairs such as *dog, cat*).

Table 1 below exemplarily lists results for MTurk-287, for which $p = \text{English}$. We notice two trends. First, for BBA, results can roughly be partitioned into three classes. The languages pt, es, fr, it have best performances as second languages with δ values between 54% and close to 60%; the next group consists of da, nl, ro, de, el, bg, sv, sl, cs with values of around 50%; finally, fi, pl, hu, lv, lt, sk, et perform worst as second languages with δ values of around 47%. So, for BBA, the choice of second language has evidently a considerable effect in that there is $\sim 26\%$ difference in performance between best second language, $\ell = \text{it}$, and worst second languages, $\ell = \text{pl/sk/et}$. Moreover, it is apparently better to choose (semantically) similar languages — with reference to the target language $p = \text{English}$ — as second language in this case. Secondly, for CCA, variation in results is much less pronounced. For example, the best second languages, et/lv, are just roughly 5.5% better than the worst second language, lt. Moreover, it is not evident, on first view, that performance results depend on language similarity in this case.⁵

	BBA	CCA		BBA	CCA
pt	56.54	57.48	sv	50.02	56.06
es	54.87	56.76	fi	47.41	56.76
fr	54.48	56.76	pl	47.12	56.47
it	59.70	57.12	cs	49.94	56.74
da	50.49	56.49	sl	52.96	57.05
nl	49.94	57.49	hu	48.84	56.46
ro	51.44	58.10	lv	47.55	58.81
de	50.08	58.24	lt	47.49	55.66
el	51.23	56.66	sk	47.22	57.17
bg	49.90	57.04	et	47.26	58.75

Table 1: Correlation coefficients $\delta = \delta_{p,\ell}$ in % on MTurk-287 for BBA and CCA methods, respectively, for various second languages ℓ . Second languages ordered by semantic similarity to $p = \text{English}$, as determined by Eq. (1); see §4.2 for specifics.

To quantify this, we systematically compute correlation coefficients τ between the correlation coefficients $\delta = \delta_{p,\ell}$ and the language distance values $D(p, \ell)$ from Eq. (1) (see §4.2 for specifics on $D(p, \ell)$). Table 2 shows that, indeed, monolingual semantic evaluation performance is consistently positively correlated with (semantic) language similarity for BBA. In contrast, for CCA, correlation is positive in eight cases and negative in six cases; moreover, coefficients are significant in only two cases. Overall, there is a strongly positive average correlation for BBA (75.75%) and a (very) weakly positive one for CCA (10.04%).

⁵As further results, we note *en passant*: CCA performed typically better than BBA, particularly in three — MEN, WS353, SimLex999 — out of our five English datasets as well as the non-English datasets. This could be due to the fact that we trained the vectors for the skip-gram model — the monolingual vectors that form the basis for CCA — on the union of Europarl and Wikipedia, while BBA used only Wikipedia as a monolingual basis. Other explanations could be the particular default hyperparameters chosen, which may have coincidentally favored CCA, or the fact that CCA uses only 1-best word alignments for projection; see Section 5 for further discussions. Moreover, in no case did we find that either BBA or CCA outperformed the purely monolingually constructed skip-gram vectors on the English evaluation task. On the one hand, this may be due to our rather small bilingual databases — containing just roughly 200K and 1,000K parallel sentences. On the other hand, while this finding is partly at odds with Faruqui and Dyer (2014), who report large improvements for bilingual word vectors over monolingual ones in some settings, it is (more) in congruence with Lu et al. (2015) and Huang et al. (2015).

$p =$		BBA	CCA
en	WS353-All21	63.75**	-6.16
	WS353-Big10	93.33**	58.33†
	MTurk287-All21	80.75***	5.11
	MTurk287-Big10	88.33**	-21.66
	MTurk771-All21	74.28***	-19.24
	MTurk771-Big10	93.33***	11.66
	SimLex999-All21	83.60***	11.57
	SimLex999-Big10	73.33*	-20.00
	MEN-All21	70.82***	-11.27
	MEN-Big10	94.99***	41.66
de	SimLex999-de-All21	60.45**	10.07
	SimLex999-de-Big10	73.33*	-31.66
it	SimLex999-it-All21	48.57*	73.83***
	SimLex999-it-Big10	61.66†	38.33
	Avg.	75.75	10.04

Table 2: Correlation τ , in %, between language similarity and monolingual semantic evaluation performance. For example, on WS353 in the Big10 setup, the more a language, say $\ell = \text{French}$, is (semantically) similar to $p = \text{English}$, the more is it likely that correlations $\delta_{p,\ell}$ are large, when word pair similarity of $p = \text{English}$ words is measured from embedding vectors that have been projected in a joint French-English semantic embedding space. More precisely, the exact correlation values are 93.33% and 58.33%, respectively, depending on whether vectors have been projected via BBA or CCA. ‘***’ means significant at the 0.1% level; ‘**’ at the 1% level, ‘*’ at the 5% level, ‘†’ at the 10% level.

	Geo	WALS	Sem
Geo		5%/45%**	40%***/65%***
WALS			23%**/62%***
Sem			

Table 3: Correlation between dist. matrices, Mantel test, All21/Big10.

4.2 Language classification (Q2)

Finally, we perform language classification on the graphs $G_\ell^{(p)}$ as indicated in §2. Since we use two different methods for inducing bilingual word embeddings, we obtain two distance matrices.⁶ Figure 3 below shows a two-dimensional representation of all 21 languages obtained from *averaging* the BBA and CCA distance matrices in the All21 setup, together with a k -means cluster assignment for $k = 6$. We note a grouping together of es, pt, fr, en, it; nl, da, de, sv; fi, et; ro, bg, el; hu, pl, cs, sk, sl; and It, Iv. In particular, {es, pt, fr, it, en} appear to form a homogeneous group with, consequently, similar semantic associations, as captured by word embeddings. Observing that fi is relatively similar to sv, which is at odds with genealogical/structural language classifications, we test another question, namely, whether the resulting semantic distance matrix is more similar to a distance matrix based on genealogical/structural relationships or to a distance matrix based on geographic relations. To this end, we determine the degree of structural similarity between two languages as the number of agreeing features (a feature is, e.g., *Number of Cases*) in the WALS⁷ database of structural properties of languages divided by the number of total features available for the language pair (Cysouw, 2013a). For geographic distance, we use the dataset from Mayer and Zignago (2011) which lists distances between countries. We make the simplifying as-

⁶For All21, these two distance matrices have a correlation of close to 70% (Mantel test), and of 73% for Big10. Hence, overall, semantic language classification results produced by either of the two methods alone — BBA or CCA — are expected to be very similar.

⁷<http://wals.info/>




Figure 3: Two-dimensional PCA (principal component analysis) projection of average distance matrices as described in the text.

sumption that, e.g., language *it* and country Italy agree, i.e., it is spoken in Italy (exclusively). Table 3 shows that geographic distance correlates better with our semantic distance calculation than does WALS structural similarity under the Mantel test measure. This may hint at an interesting result: since semantics is changing fast, it may be more directly influenced by contact phenomena than by genealogical processes that operate on a much slower time-scale.

Note that our results are in accordance with the assumption that the probability of borrowing and geographical distance are inversely correlated (Cysouw, 2013b). In our case, this may relate to semantic loans (adopting the semantic neighborhoods of loaned words within the target language) rather than to structural or grammatical borrowings. That is, geographically related languages exhibit a higher probability to borrow words from each other together with the same range of semantic associations. At least, this hypothesis is not falsified by our experiment.

5 Discussion

Our initial expectation was that ‘distant’ second languages ℓ — in terms of language similarity — would greatly deteriorate monolingual semantic evaluations in a target language p , as we believed they would invoke ‘unusual’ semantic associations from the perspective of p . Such a finding would have been a word of caution regarding with which language to embed a target language p in a joint semantic space, if this happens for the sake of improving monolingual semantic similarity in p . We were surprised to find that only BBA was sensitive to language similarity in our experiments in this respect, whereas CCA seems quite robust against choice of second language. An explanation for this finding may be the different manners in which both methods induce joint embedding spaces: While CCA takes independently constructed vector spaces of two languages, BBA jointly optimizes mono- and bilingual constraints and may thus be more sensitive to the interplay, and relation, between both languages. Another plausible explanation is that CCA uses only *1-best* alignments for projecting two languages in a joint semantic space. Thus, it may be less sensitive to varying polysemous associations across different languages (cf. our *vir* example in Section 1), and hence less adequate for capturing cross-lingual polysemy.⁸

In terms of language similarity, we mention that our approach is formally similar to approaches as in

⁸Thus, we would also expect CCA to perform better in monolingual intrinsic evaluations (as our experiments have partly confirmed) and BBA to perform better in multilingual intrinsic evaluations. We thank one reviewer for pointing this out.

(Eger et al., 2015; Asgari and Mofrad, 2016) and others. Namely, we construct graphs, one for each language, and compare them to determine language distance. Compared to Eger et al. (2015), our approach differs in that they use translations in a second language ℓ to measure similarity between pivot language p words. This idea also underlies very well-known lexical semantic resources such as the paraphrase database (PPDB) (Bannard and Callison-Burch, 2005; Ganitkevitch et al., 2013); see also Eger and Se-jane (2010). In contrast, we directly use bilingual embeddings for this similarity measurement by jointly embedding p and ℓ , which are arguably best suited for this task. Our approach also differs from Eger et al. (2015) in that we do not apply a random-surfer process to our semantic graphs.

We finally note that the linguistic problem of (semantic) language classification, as we consider, involves some vagueness as there is de facto no gold standard that we can compare to. Reasonably, however, languages should be semantically similar to a degree that reflects structural, genealogical, and contact relationships. One approach may then be to disentangle or, as we pursued here, (relatively) weigh each of these effects.

From an application perspective, our approach allows for enriching (automatically generated) lexica. This relates, for example, to the generation of sentiment lexica listing prior polarities for selected lexemes (Sonntag and Stede, 2015). Since the acquisition of such specialized information (e.g., by annotation) is cost-intensive, approaches are needed that allow for automatically generating or extending such resources especially in the case of historical languages (e.g., Latin). Here our idea is to start from pairs of semantically (most) similar languages in order to induce polarity cues for words in the target language as a function of their distances to selected seed words in the pivot language, for which polarities are already known. By averaging over groups of semantically related pivot languages, for which sentiment lexica already exist, the priority listings for the desired target language may stabilize. Obviously, this procedure can be applied to whatever lexical information to be annotated automatically (e.g., grammatical or semantic categories like agency, animacy etc. as needed, for example, for semantic role labeling (Palmer et al., 2010)).

A second application scenario relates to measuring (dis-)similarities of translations and their source texts (Baker, 1993): starting from our model of bilingual semantic spaces, we may ask, for example, whether words for which several alternatives exist within the target language tend to be translated by candidates that retain most of their associations within the source language – possibly in contradiction to frequency effects. Such a finding would be in line with Toury’s notion of interference (Toury, 1995) according to which translations reflect characteristics of the source language – the latter leaves, so to speak, fingerprints within the former. Such a finding would bridge between the notion of interference in translation studies and distributional semantics based on deep learning.

6 Related work

Besides the mono- and multilingual word vector representation research that forms the basis of our work and which has already been referred to, we mention the following three related approaches to language classification. Koehn (2005) compares down-stream task performance in SMT to language family relationship, finding positive correlation. Cooper (2008) measures semantic language distance via bilingual dictionaries, finding that French appears to be semantically closer to Basque than to German, supporting our arguments on contact as co-determining semantic language similarity. Bamman et al. (2014) and Kulkarni et al. (2015b) study semantic distance between dialects of English by comparing region specific word embeddings.

Studying geographic variation of (different) languages is also closely related to studying temporal variation within one and the same language (Kulkarni et al., 2015a), with one crucial difference being the need to find a common representation in the former case. Word embeddings — in particular, monolingual ones — can also be used to address the latter scenario (Eger and Mehler, 2016; Hamilton et al., 2016).

In terms of classifying languages, the work that is closest to ours is that of Asgari and Mofrad (2016). A key difference between their approach and ours is that, in order to achieve a common representation between languages, they translate words. This has the disadvantage that translation pairs need to be known, which typically requires large amounts of parallel text. In contrast, bilingual word embeddings, which

form the basis of our experiments, can be generated from as few as ten translation pairs, as demonstrated in Zhang et al. (2016).

There is by now a long-standing tradition that compares languages via analysis of complex networks that encode their words and the (semantic) relationships between them (Cancho and Solé, 2001; Gao et al., 2014). These studies often only look at very abstract statistics of networks such as average path lengths and clustering coefficients, rather than analyzing them on a level of content of their nodes and edges. In addition, they often substitute co-occurrence as a proxy for semantic similarity. However, as Asgari and Mofrad (2016) point out, co-occurrence is a naive estimate of similarity; e.g., synonyms rarely co-occur.

7 Conclusion

Using English, German and Italian as pivot languages, we show that the choice of the second language may significantly matter when the resulting space is used for monolingual semantic evaluation tasks. More specifically, we show that the goodness of this choice is influenced by genealogical similarity and by (geographical) language contact. This finding may be important for the question which languages to integrate in multilingual embedding spaces (Huang et al., 2015). Moreover, we show that semantic language similarity — estimated on the basis of bilingual embedding spaces as suggested in this work — may be better predicted by contact than by genealogical relatedness. The validation of this hypothesis by means of bigger data sets will be the object of future work.

References

- Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Polyglot: Distributed word representations for multilingual nlp. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 183–192, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Ehsaneddin Asgari and Mohammad R.K. Mofrad. 2016. Comparing fifty natural languages and twelve genetic languages using word embedding language divergence (weld) as a quantitative measure of language distance. In *Proceedings of the Workshop on Multilingual and Cross-lingual Methods in NLP*, pages 65–74, San Diego, California, June. Association for Computational Linguistics.
- Mona Baker. 1993. Corpus linguistics and translation studies: Implications and applications. In Mona Baker, Francis Gill, and Elena Tognini-Bonelli, editors, *Text and Technology: In Honour of John Sinclair*, pages 233–250, Amsterdam/Philadelphia. John Benjamins.
- David Bamman, Chris Dyer, and Noah A. Smith. 2014. Distributed representations of geographically situated language. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 828–834, Baltimore, Maryland, June. Association for Computational Linguistics.
- Colin Bannard and Chris Callison-Burch. 2005. Paraphrasing with bilingual parallel corpora. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, ACL '05*, pages 597–604, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 238–247, Baltimore, Maryland, June. Association for Computational Linguistics.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *J. Mach. Learn. Res.*, 3:1137–1155, March.
- Elia Bruni, Nam Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *J. Artif. Int. Res.*, 49(1):1–47, January.
- Horst Bunke and Kim Shearer. 1998. A graph distance metric based on the maximal common subgraph. *Pattern Recognition Letters*, 19(3-4):255–259, March.
- Ramon F. Cancho and Richard V. Solé. 2001. The small world of human language. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 268(1482):2261–2265, November.

- A. P. 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 *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1853–1861.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 160–167, New York, NY, USA. ACM.
- Martin C. Cooper. 2008. Measuring the semantic distance between languages from a statistical analysis of bilingual dictionaries. *Journal of Quantitative Linguistics*, 15(1):1–33.
- Michael Cysouw, 2013a. *Approaches to Measuring Linguistic Differences*, chapter Predicting language learning difficulty, pages 57–82. De Gruyter Mouton.
- Michael Cysouw. 2013b. Disentangling geography from genealogy. In Peter Auer, Martin Hilpert, Anja Stukenbrock, and Benedikt Szmezcanyi, editors, *Space in Language and Linguistics: Geographical, Interactional, and Cognitive Perspectives*, pages 21–37, Berlin. De Gruyter Mouton.
- Chris Dyer, Adam Lopez, Juri Ganitkevitch, Johnathan Weese, Ferhan Ture, Phil Blunsom, Hendra Setiawan, Vladimir Eidelman, and Philip Resnik. 2010. cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Steffen Eger and Alexander Mehler. 2016. On the linearity of semantic change: Investigating meaning variation via dynamic graph models. In *Proceedings of ACL 2016*. Association for Computational Linguistics.
- Steffen Eger and Ineta Sejane. 2010. Computing Semantic Similarity from Bilingual Dictionaries. In *Proceedings of the 10th International Conference on the Statistical Analysis of Textual Data (JADT-2010)*, pages 1217–1225, Rome, Italy, June. JADT-2010.
- Steffen Eger, Niko Schenk, and Alexander Mehler. 2015. Towards semantic language classification: Inducing and clustering semantic association networks from Europarl. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics*, pages 127–136, Denver, Colorado, June. Association for Computational Linguistics.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of EACL*.
- Manaal Faruqui, Jesse Dodge, Sujay K. Jauhar, Chris Dyer, Eduard Hovy, and Noah A. Smith. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of NAACL*.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppín. 2002. Placing search in context: The concept revisited. *ACM Transactions on Information Systems*, 20(1):116–131, January.
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. PPDB: The paraphrase database. In *Proceedings of NAACL-HLT*, pages 758–764, Atlanta, Georgia, June. Association for Computational Linguistics.
- Yuyang Gao, Wei liang, Yuming Shi, and Qiuling Huang. 2014. Comparison of directed and weighted co-occurrence networks of six languages. *Physica A: Statistical Mechanics and its Applications*, (393):579–589.
- Stephan Gouws and Anders Søgaard. 2015. Simple task-specific bilingual word embeddings. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1386–1390, Denver, Colorado, May–June. 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 Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 748–756.
- William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.

- Karl Moritz Hermann and Phil Blunsom. 2014. Multilingual models for compositional distributed semantics. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 58–68.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. 2012. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers - Volume 1, ACL '12*, pages 873–882, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kejun Huang, Matt Gardner, Evangelos E. Papalexakis, Christos Faloutsos, Nikos D. Sidiropoulos, Tom M. Mitchell, Partha Pratim Talukdar, and Xiao Fu. 2015. Translation invariant word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1084–1088.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. 2012. Inducing crosslingual distributed representations of words. In *Proceedings of COLING 2012*, pages 1459–1474, Mumbai, India, December. The COLING 2012 Organizing Committee.
- Philipp Koehn. 2005. Europarl: A Parallel Corpus for Statistical Machine Translation. In *Conference Proceedings: the tenth Machine Translation Summit*, pages 79–86, Phuket, Thailand. AAMT, AAMT.
- Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015a. Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web, WWW '15*, pages 625–635, New York, NY, USA. ACM.
- Vivek Kulkarni, Bryan Perozzi, and Steven Skiena. 2015b. Freshman or fresher? quantifying the geographic variation of internet language. *CoRR*, abs/1510.06786.
- Ira Leviant and Roi Reichart. 2015. Separated by an un-common language: Towards judgment language informed vector space modeling. *CoRR*, abs/1508.00106.
- Ang Lu, Weiran Wang, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. Deep multilingual correlation for improved word embeddings. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 250–256, Denver, Colorado, May–June. Association for Computational Linguistics.
- Thierry Mayer and Soledad Zignago. 2011. Notes on ceptis distances measures: The geodist database. Working Papers 2011-25, CEPII.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. 2014. Efficient non-parametric estimation of multiple embeddings per word in vector space. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1059–1069, Doha, Qatar, October. Association for Computational Linguistics.
- Martha Palmer, Daniel Gildea, and Nianwen Xue. 2010. *Semantic Role Labeling*. Morgan & Claypool Publishers, San Rafael.
- 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, Doha, Qatar, October. Association for Computational Linguistics.
- 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, WWW '11*, pages 337–346, New York, NY, USA. ACM.
- Sascha Rothe and Hinrich Schütze. 2014. Cosimrank: A flexible & efficient graph-theoretic similarity measure. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1392–1402, Baltimore, Maryland, June. Association for Computational Linguistics.

- Jonathan Sonntag and Manfred Stede. 2015. Sentiment analysis: Whats your opinion? In Chris Biemann and Alexander Mehler, editors, *Text Mining: From Ontology Learning to Automated Text Processing Applications. Festschrift in Honor of Gerhard Heyer*, Theory and Applications of Natural Language Processing. Springer, Heidelberg.
- Gideon Toury. 1995. *Descriptive Translation Studies and Beyond*. John Benjamins, Amsterdam/Philadelphia.
- Simon Šuster, Ivan Titov, and Gertjan van Noord. 2016. Bilingual learning of multi-sense embeddings with discrete autoencoders. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1346–1356, San Diego, California, June. Association for Computational Linguistics.
- Yuan Zhang, David Gaddy, Regina Barzilay, and Tommi Jaakkola. 2016. Ten pairs to tag – multilingual pos tagging via coarse mapping between embeddings. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1307–1317, San Diego, California, June. Association for Computational Linguistics.