

Cross-Lingual Alignment of Contextual Word Embeddings, with Applications to Zero-shot Dependency Parsing

Tal Schuster^{*,1}, Ori Ram^{*,2}, Regina Barzilay¹, Amir Globerson²

¹Computer Science and Artificial Intelligence Lab, MIT

²Tel Aviv University

{tals, regina}@csail.mit.edu, {ori.ram, gamir}@cs.tau.ac.il

Abstract

We introduce a novel method for multilingual transfer that utilizes deep contextual embeddings, pretrained in an unsupervised fashion. While contextual embeddings have been shown to yield richer representations of meaning compared to their static counterparts, aligning them poses a challenge due to their dynamic nature. To this end, we construct context-independent variants of the original monolingual spaces and utilize their mapping to derive an alignment for the context-dependent spaces. This mapping readily supports processing of a target language, improving transfer by context-aware embeddings. Our experimental results demonstrate the effectiveness of this approach for zero-shot and few-shot learning of dependency parsing. Specifically, our method consistently outperforms the previous state-of-the-art on 6 tested languages, yielding an improvement of 6.8 LAS points on average.¹

1 Introduction

Multilingual embedding spaces have been demonstrated to be a promising means for enabling cross-lingual transfer in many natural language processing tasks (e.g. Ammar et al. (2016); Lample et al. (2018)). Similar to how universal part-of-speech tags enabled parsing transfer across languages (Petrov et al., 2012), multilingual word embeddings further improve transfer capacity by enriching models with lexical information. Since this lexical representation is learned in an unsupervised fashion and thus can leverage large amounts of raw data, it can capture a more nuanced representation of meaning than unlexicalized transfer. Naturally, this enrichment is trans-

lated into improved transfer accuracy, especially in low-resource scenarios (Guo et al., 2015).

In this paper, we are moving further along this line and exploring the use of contextual word embeddings for multilingual transfer. By dynamically linking words to their various contexts, these embeddings provide a richer semantic and syntactic representation than traditional context-independent word embeddings (Peters et al., 2018). A straightforward way to utilize this richer representation is to directly apply existing transfer algorithms on the contextual embeddings instead of their static counterparts. In this case, however, each token pair is represented by many different vectors corresponding to its specific context. Even when supervision is available in the form of a dictionary, it is still unclear how to utilize this information for multiple contextual embeddings that correspond to a word translation pair.

In this paper, we propose a simple but effective mechanism for constructing a multilingual space of contextual embeddings. Instead of learning the alignment in the original, complex contextual space, we drive the mapping process using context-independent embedding anchors. We obtain these anchors by factorizing the contextual embedding space into context-independent and context-dependent parts. Operating at the anchor level not only compresses the space, but also enables us to utilize a word-level bilingual dictionary as a source of supervision, if available. Once the anchor-level alignment is learned, it can be readily applied to map the original spaces with contextual embeddings.

Clearly, the value of word embeddings depends on their quality, which is determined by the amount of raw data available for their training (Jiang et al., 2018). We are interested in expanding the above approach to the truly low-resource

* Equal contribution

¹Code and models: <https://github.com/TalSchuster/CrossLingualELMo>.

scenario, where a language not only lacks annotations, but also has limited amounts of raw data. In this case, we can also rely on a data rich language to stabilize monolingual embeddings of the resource-limited language. As above, context-independent anchors are informing this process. Specifically, we introduce an alignment component to the loss function of the language model, pushing the anchors to be closer in the joint space. While this augmentation is performed on the static anchors, the benefit extends to the contextual embeddings space in which we operate.

We evaluate our aligned contextual embeddings on the task of zero-shot cross-lingual dependency parsing. Our model consistently outperforms previous transfer methods, yielding absolute improvement of 6.8 LAS points over the prior state-of-the-art (Ammar et al., 2016). We also perform comprehensive studies of simplified variants of our model. Even without POS tag labeling or a dictionary, our model performs on par with context-independent models that do use such information. Our results also demonstrate the benefits of this approach for few-shot learning, i.e. processing languages with limited data. Specifically, on the Kazakh tree-bank from the recent CoNLL 2018 shared task with only 38 trees for training, the model yields 5 LAS points gain over the top result (Smith et al., 2018a).

2 Related work

Multilingual Embeddings The topic of cross-lingual embedding alignment is an active area of research (Mikolov et al., 2013; Xing et al., 2015; Dinu and Baroni, 2014; Lazaridou et al., 2015; Zhang et al., 2017). Our work most closely relates to MUSE (Conneau et al., 2018a), which constructs a multilingual space by aligning monolingual embedding spaces. When a bilingual dictionary is provided, their approach is similar to those of (Smith et al., 2017; Artetxe et al., 2017). MUSE extends these methods to the unsupervised case by constructing a synthetic dictionary. The resulting alignment achieves strong performance in a range of NLP tasks, from sequence labeling (Lin et al., 2018) to natural language inference (Conneau et al., 2018b) and machine translation (Lample et al., 2018; Qi et al., 2018). Recent work further improves the performance on both the supervised (Joulin et al., 2018) and unsupervised (Grave et al., 2018b; Alvarez-Melis and

Jaakkola, 2018; Hoshen and Wolf, 2018) settings for context-independent embeddings.

While MUSE operates over token based embeddings, we are interested in aligning contextual embeddings, which have shown their benefits in several monolingual applications (Peters et al., 2018; McCann et al., 2017; Howard and Ruder, 2018; Radford et al., 2018; Devlin et al., 2018). However, this expansion introduces new challenges which we address in this paper.

In a concurrent study, Aldarmaki and Diab (2019) introduced an alignment that is based only on word pairs in the same context, using parallel sentences. Our method achieves better word translations without relying on such supervision.

Our work also relates to prior approaches that utilize bilingual dictionaries to improve embeddings that were trained on small datasets. For instance, Xiao and Guo (2014) represent word pairs as a mutual vector, while Adams et al. (2017) jointly train cross-lingual word embeddings by replacing the predicted word with its translation. To utilize a dictionary in the contextualized case, we include a soft constraint that pushes those translations to be similar in their context-independent representation. A similar style of regularization was shown to be effective for cross-domain transfer of word embeddings (Yang et al., 2017).

Multilingual Parsing In early work on multilingual parsing, transfer was commonly implemented using delexicalized representation such as part-of-speech tags (McDonald et al., 2011; Petrov et al., 2012; Naseem et al., 2012; Tiedemann, 2015).

Another approach for cross-lingual parsing includes annotation projection and treebank translation (Xiao and Guo, 2015; Wang and Eisner, 2016; Tiedemann, 2017), which mostly require some source of supervision.

Advancements in multilingual word representations opened a possibility of lexicalized transfer. Some of these approaches start by aligning monolingual embedding spaces (Zhang and Barzilay, 2015; Guo et al., 2015, 2016; Ammar et al., 2016), and using resulting word embeddings as word representations instead of universal tags. Other approaches are learning customized multilingual syntactic embeddings bootstrapping from universal POS tags (Duong et al., 2015). While some models also learn a language embedding (Ammar et al., 2016; de Lhoneux et al., 2018), it is unfeasible in a zero-shot scenario.

$D_{cos}(\bar{e}_i, \bar{e}_j)$	$D_{cos}(\bar{e}_i, e_{i,c})$	
	ALL WORDS	HOMONYMS
0.85 (± 0.09)	0.18 (± 0.04)	0.21 (± 0.04)

Table 1: Average cosine distances between pairs of embedding anchors (left column) and between contextualized embeddings of words to their corresponding anchor. The right column includes these distances only for homonyms, whereas the center column is averaged across all words. Only alphabetic words with at least 100 occurrences were included.

In all of the above cases, token-level embeddings are used. Inspired by strong results of using contextualized embeddings in monolingual parsing (Che et al., 2018; Wang et al., 2018; Clark et al., 2018), we aim to utilize them in the multilingual transfer case. Our results demonstrate that richer representation of lexical space does lead to significant performance gains.

3 Aligning Contextual Word Embeddings

In this section we describe several approaches for aligning context-dependent embeddings from a source language s to a target language t . We address multiple scenarios, where different amounts of supervision and data are present. Our approach is motivated by interesting properties of context-dependent embeddings, which we discuss later.

We begin with some notations:

- **Context Dependent Embeddings:** Given a context c and a token i , we denote the embedding of i in the context c by $e_{i,c}$. We use $e_{i,\cdot}$ to denote the point cloud of all contextual embeddings for token i .
- **Embedding Anchor:** Given a token i we denote the anchor of its context dependent embeddings by \bar{e}_i , where:

$$\bar{e}_i = \mathbb{E}_c [e_{i,c}] . \quad (1)$$

In practice, we calculate the average over a subset of the available unlabeled data.

- **Shift From Mean:** For any embedding $e_{i,c}$ we can therefore define the shift $\hat{e}_{i,c}$ from the average via:

$$e_{i,c} = \bar{e}_i + \hat{e}_{i,c} . \quad (2)$$




Figure 1: A two dimensional PCA showing examples of contextual representations for four Spanish words. Their corresponding anchors are presented as a star in the same color. (best viewed in color)

- **Embedding Alignment:** Given an embedding $e_{i,c}^s$ in s , we want to generate an embedding $e_{i,c}^{s \rightarrow t}$ in the target language space, using a linear mapping $W^{s \rightarrow t}$. Formally, our alignment is always of the following form:

$$e_{i,c}^{s \rightarrow t} = W^{s \rightarrow t} e_{i,c}^s . \quad (3)$$

3.1 The Geometry of Context-Dependent Embeddings

A given token i can generate multiple vectors $e_{i,c}$, each corresponding to a different context c . A key question is how the point cloud $e_{i,\cdot}$ is distributed. In what follows we explore this structure, and reach several conclusions that will motivate our alignment approach. The following experiments are performed on ELMo (Peters et al., 2018).

Point Clouds are Well Separated A cloud $e_{i,\cdot}$ corresponds to occurrences of the word i in different contexts. Intuitively, we would expect its points to be closer to each other than to points from $e_{j,\cdot}$ for a different word j . Indeed, when measuring similarity between points $e_{i,c}$ and their anchor \bar{e}_i , we find that these are much more similar than anchors of different words \bar{e}_i and \bar{e}_j (see Table 1). This observation supports our hypothesis that anchor-driven alignment can guide the construction of the alignment for the contextual space. A visualized example of the contextualized representations of four words is given in Figure 1, demonstrating the appropriateness of their anchors. Still, as previous studies have shown, and




Figure 2: Contextual embeddings for the English word “bear” and its two possible translations in Spanish — “oso” (animal) in blue and “tener” (to have) in red. The figure shows a two dimensional PCA for the aligned space of the two languages. The symbols are the anchors, the clouds represent the distribution of the contextualized Spanish words, and the black dots are for contextualized embeddings of “bear”. The gray colored triangles show the anchors of the English words “dog”, “elephant”, “cat”, from left to right respectively.

as our results point, the context component is very useful for downstream tasks.

Homonym Point Clouds are Multi-Modal

When a word i has multiple distinct senses, we might expect the embeddings for i to reflect this by separating into multiple distinct clouds, one for each meaning. Figure 2 demonstrates that this indeed happens for the English word “bear”. Furthermore, it can be seen that after alignment (Section 3.3) with Spanish, the distinct point clouds are aligned with their corresponding distinct words in Spanish. See App. D for another example.

We examined the shift from mean for a list of 250 English homonyms from Wikipedia.² As Table 1 shows, the shift of these words is indeed slightly higher than it is for other words. However, they still remain relatively close to their per-token anchor. Therefore, these anchors can still serve as a good approximation for learning alignments.

3.2 Context-Independent Alignment

We begin by briefly reviewing previous approaches for aligning context-independent embeddings, as they are generalized in this work to the contextual case. We denote the embedding of a word i by e_i . At first, assume we are given word pairs $\{(e_i^s, e_i^t)\}$ from a source language s and a

²https://en.wikipedia.org/wiki/List_of_true_homonyms

target language t , and we look for a mapping between those. Mikolov et al. (2013) proposed to learn a linear transformation whereby e_i^s is approximated via $W e_i^t$, for a learned matrix W . We focus on methods that follow this linear alignment. The alignment matrix is found by solving:

$$W^{s \rightarrow t} = \operatorname{argmin}_{W \in O_d(\mathbb{R})} \sum_{i=1}^n \|W e_i^s - e_i^t\|^2, \quad (4)$$

where $O_d(\mathbb{R})$ is the space of orthogonal matrices. This constraint was proposed by Xing et al. (2015) in order to preserve inter-lingual relations. Under this constraint, Eq. 4 is an instance of the orthogonal Procrustes problem, which has a closed-form solution $W^{s \rightarrow t} = UV^T$. The columns of U and V are the left and right singular vectors of the multiplication of the source and (transposed) target embedding matrices.

For the *unsupervised* case (i.e. when a dictionary is absent), Conneau et al. (2018a) (MUSE) suggested to learn the alignment via adversarial training, such that a discriminator is trained to distinguish between target and aligned source embeddings. Thereafter, a *refinement* procedure is applied iteratively as follows. First, a dictionary is built dynamically using the current alignment such that only words with high confidence are considered. Using the dictionary, the alignment matrix is re-calculated as in the supervised case.

3.3 Context-Dependent Alignment

We next turn our attention to the main task of this paper, which is aligning context-dependent embeddings. We now describe our generalization of the methods described in Section 3.2 for this case. The first two methods are based only on anchors while the third one uses the contextual vectors themselves. Altogether, we suggest three alignment procedures, one aimed for the supervised and two for the unsupervised cases.

Supervised Anchored Alignment As a first step, we are assuming access to a dictionary for the source and target domains. For each source word i denote by $D(i)$ the corresponding word in the target language.³

In the context-dependent case, Eq. 4 is no longer well-defined, as there are many corresponding vectors to both the source and the target

³In practice, we may have multiple target words for a single source word, and the extension is straight-forward.

words. However, this challenge can be addressed by aligning the vectors \bar{e}_i for which we do have one per word. This is motivated by our observations in Section 3.1 that context-dependent embeddings are well clustered around their centers.

Thus, in the case where a dictionary is available, we solve Eq. 4 with token anchors as inputs.

We emphasize that by constraining $W^{s \rightarrow t}$ to be orthogonal, we also preserve relations between $\hat{e}_{i,c}$ and $\hat{e}_{i,c'}$ that represent the contextual information.

Unsupervised Anchored Alignment In this setting, no dictionary is present. As in the supervised case, we can naturally extend a context-independent alignment procedure to the contextual space by leveraging the anchor space \bar{e}_i . This can be done using the adversarial MUSE framework proposed by [Conneau et al. \(2018a\)](#) and described at the end of Section 3.2.

Unsupervised Context-based Alignment Alternatively, the alignment could be learned directly on the contextual space. To this end, we follow again the adversarial algorithm of MUSE, but for each word we use multiple embeddings induced by different contexts, rather than the word anchor.

This context-based alignment presents opportunities but also introduces certain challenges. On the one hand, it allows to directly handle homonyms during the training process. However, empirically we found that training in this setting is less stable than unsupervised anchored alignments.

Refinement As a final step, for both of the unsupervised methods, we perform the refinement procedure that is incorporated in MUSE (end of Section 3.2). In order to synthesize a dictionary, we use distance in the anchor space.

3.4 Learning Anchored Language Models

Thus far we assumed that embeddings for both source and target languages are pretrained separately. Afterwards, the source is mapped to the target in a second step via a learned mapping. However, this approach may not work well when raw data for the source languages is scarce, resulting in deficient embeddings. In what follows, we show how to address this problem when a dictionary is available. We focus on embeddings that are learned using a language model objective but

this can be easily generalized to other objectives as well.

Our key idea is to constrain the embeddings across languages such that word translations will be close to each other in the embedding space. This can serve as a regularizer for the resource-limited language model. In this case, the anchors are the model representations prior to its context-aware components (e.g., the inputs to ELMo’s LSTM).

Denote the anchor for word i in language s by v_i^s . Now, assume we have trained a model for the target language and similarly have embeddings v_i^t . We propose to train the source model with an added regularization term as follows:

$$\lambda_{\text{anchor}} \cdot \sum_i \|v_i^s - v_{D(i)}^t\|_2^2, \quad (5)$$

where λ_{anchor} is a hyperparameter. This regularization has two positive effects. First, it reduces overfitting by reducing the effective number of parameters the model fits (e.g., if the regularizer has large coefficient, these parameters are essentially fixed). Second, it provides a certain level of alignment between the source and target language since they are encouraged to use similar anchors.

4 Multilingual Dependency Parsing

Now that we presented our method for aligning contextual embeddings, we turn to evaluate it on the task of cross-lingual dependency parsing. We first describe our baseline model, and then show how our alignment can easily be incorporated into this architecture to obtain a multilingual parser.

Baseline Parser Most previous cross-lingual dependency parsing models used transition-based models ([Ammar et al., 2016](#); [Guo et al., 2016](#)). We follow [Che et al. \(2018\)](#); [Wang et al. \(2018\)](#); [Clark et al. \(2018\)](#) and use a first-order graph-based model. Specifically, we adopt the neural edge-scoring architecture from [Dozat and Manning \(2017\)](#); [Dozat et al. \(2017\)](#), which is based on [Kiperwasser and Goldberg \(2016\)](#). We now briefly review this architecture. Given a sentence s , let e_i and p_i be its word and POS-tag embeddings. These are concatenated and fed into a Bi-LSTM to produce token-level contextual representations r_i . Four Multi-Layer Perceptrons are applied on these vectors, resulting in new representations $h_i^{\text{arc-dep}}$, $h_i^{\text{arc-head}}$, $h_i^{\text{rel-dep}}$ and $h_i^{\text{rel-head}}$

for each word i . Arc scores are then obtained by:

$$s_{ij}^{arc} = \left(\mathbf{h}_i^{arc-head} \right)^T \left(U^{arc} \mathbf{h}_j^{arc-dep} + \mathbf{b}^{arc} \right). \quad (6)$$

Additionally, the score for predicting the dependency label r for an edge (i, j) is defined as

$$s_{(i,j),r}^{rel} = \left(\mathbf{h}_i^{rel-head} \right)^T U_r^{rel} \mathbf{h}_j^{rel-dep} + \left(\mathbf{u}_r^{rel-head} \right)^T \mathbf{h}_i^{rel-head} + \left(\mathbf{u}_r^{rel-dep} \right)^T \mathbf{h}_j^{rel-dep} + b_r. \quad (7)$$

At test time, MST is calculated to ensure valid outputs.

Multilingual Parsing with Alignment We now extend this model, in order to effectively use it for transfer learning. First, we include contextualized word embeddings by replacing the static embeddings with a pre-trained ELMo (Peters et al., 2018) model (instead of e_i). Second, we share all model parameters across languages and use the contextual word embeddings after they are aligned to a joint space J . Formally, if s is a sentence of language ℓ , contextual word embeddings are obtained via:

$$\mathbf{e}_{i,s}^{\ell \rightarrow J} = W^{\ell \rightarrow J} \mathbf{e}_{i,s}, \quad (8)$$

where $W^{\ell \rightarrow J}$ is the alignment matrix from language ℓ to the joint space.⁴ This alignment is learned apriori and kept fixed during parser training. This setup is applicable for both single and multiple training languages. For the tested language, training data could be available, sometimes limited (few-shot), or absent (zero-shot). The alignment methods are described in detail in Section 3.

In their paper, Peters et al. (2018) suggest to output a linear combination over the representations of each layer of ELMo, learning these weights jointly with a downstream task. Our alignment is learned separately for each layer. Therefore, we keep the weights of the combination fixed during the training to ensure that the parser’s inputs are from the joint cross-lingual space. Alternatively, one can share the weights of the combination between the languages and learn them.

All the above modifications are at the word embedding level, making them applicable to any other NLP model that uses word embeddings.

⁴We use the space of the training language as our joint space and align the tested language to it. In the multi-source scenario, we align all embeddings to English.

5 Experimental Setup

Contextual Embeddings We use the ELMo model (Peters et al., 2018) with its default parameters to generate embeddings of dimension 1024 for all languages. For each language, training data comprises Wikipedia dumps⁵ that were tokenized using UDpipe (Straka and Straková, 2017). We randomly shuffle the sentences and, following the setting of ELMo, use 95% of them for training and 5% for evaluation.

Alignment We utilize the MUSE framework⁶ (Conneau et al., 2018a) and the dictionary tables provided by them. The \bar{e}_i (anchor) vectors for the alignment are generated by computing the average of representations on the evaluation set (except for the limited unlabeled data case). To evaluate our alignment, we use the anchors to produce word translations. For all experiments we use the 50k most common words in each language.

Dependency Parsing We used the biaffine parser implemented in AllenNLP (Gardner et al., 2018), refactored to handle our modifications as described in Section 4.⁷ The parser is trained on trees from a single or multiple languages, as described in each setting (Section 6). For the multiple case, we randomly alternate between the available languages, i.e. at each iteration we randomly choose one language and sample a corresponding batch. Dropout (Srivastava et al., 2014) is applied on ELMo representations, Bi-LSTM representations and outputs of MLP layers. We also apply early stopping, where validation accuracy is measured as average LAS score on the development set across all training languages. The parser hyperparameters are the same as Dozat et al. (2017) except we reduce the POS tag embedding size from 100 to 50 and increase the head/dependent MLP dimension from 400 to 500. All hyperparameter values used are listed in App. C.

From experiments on the English tree-bank, we found that using the outputs of the first LSTM layer is as good as learning a combination.⁸ This

⁵<https://lindat.mff.cuni.cz/repository/xmlui/handle/11234/1-1989>

⁶<https://github.com/facebookresearch/MUSE/>

⁷<https://github.com/TalSchuster/allennlp-MultiLang>

⁸This was concurrently justified by Liu et al. (2019), showing that the first layer alone can perform better than a mixture.

ALIGNMENT METHOD	DE	ES	FR	IT	PT	SV	AVERAGE
SUPERVISED ANCHORED	78	85	86	82	74	68	79
UNSUPERVISED ANCHORED	63	61	70	58	35	22	52
+ REFINE	72	74	81	77	53	33	65
UNSUPERVISED CONTEXT-BASED	57	68	59	57	53	*	49
+ REFINE	73	82	77	73	66	*	62

Table 2: Word translation to English precision @5 using CSLS (Conneau et al., 2018a) with a dictionary (supervised) and without (unsupervised) for German (DE), Spanish (ES), French (FR), Italian (IT), Portuguese (PT) and Swedish (SV). Each of the unsupervised results is followed by a line with the results post the anchor-based refinement steps. * stands for 'Failed to converge'.

MODEL	DE	ES	FR	IT	PT	SV	AVERAGE
Zhang and Barzilay (2015)	54.1	68.3	68.8	69.4	72.5	62.5	65.9
Guo et al. (2016)	55.9	73.1	71.0	71.2	78.6	69.5	69.9
Ammar et al. (2016)	57.1	74.6	73.9	72.5	77.0	68.1	70.5
ALIGNED FASTTEXT	61.5	78.2	76.9	76.5	83.0	70.1	74.4
ALIGNED \bar{e}	58.0	76.7	76.7	76.1	79.2	71.9	73.1
OURS	65.2	80.0	80.8	79.8	82.7	75.4	77.3
OURS, NO DICTIONARY	64.1	77.8	79.8	79.7	79.1	69.6	75.0
OURS, NO POS	61.4	77.5	77.0	77.6	73.9	71.0	73.1
OURS, NO DICTIONARY, NO POS	61.7	76.6	76.3	77.1	69.1	54.2	69.2

Table 3: Zero-shot cross lingual LAS scores compared to previous methods, for German (DE), Spanish (ES), French (FR), Italian (IT), Portuguese (PT) and Swedish (SV). Aligned FASTTEXT and \bar{e} context-independent embeddings are also presented as baselines. The bottom three rows are models that don't use POS tags at all and/or use an unsupervised anchored alignment. Corresponding UAS results are provided in App. B.

agrees with Belinkov et al. (2017), showing that lower layers capture more syntactic information. Therefore, we fix the weights over ELMo layers to $[0, 1, 0]$, i.e. using only representations from the first LSTM layer.

Evaluation Scenarios for Dependency Parsing

For a fair comparison, we use the same setting as used by previous models for each scenario. Our main model (which we refer to as OURS) is using a SUPERVISED ANCHORED ALIGNMENT (Section 3.3) to align the multilingual pretrained ELMo embeddings which are used by the parser. We compare against several variants of our model:

- ALIGNED FASTTEXT: instead of ELMo, we use FASTTEXT pretrained embeddings (Grave et al., 2018a), aligned to English using MUSE.
- ALIGNED \bar{e} : instead of contextualized embeddings, we use the anchors themselves as fixed embeddings, aligned to English.

- NO DICTIONARY: we assume the absence of a dictionary and use UNSUPERVISED ANCHORED ALIGNMENT.

- NO POS: no use of part of speech tags.

6 Results

Alignment As mentioned above, we use outputs of the first LSTM layer of ELMo in our parsing experiments. Therefore, we present the alignment accuracy for those in Table 2, summarizing the precision@5 word-translation from 6 languages to English. Results for the other layers are presented in App. A. As expected, supervised alignments outperform unsupervised ones by a large margin. Between the two unsupervised methods, the context-based alignment achieved significantly better results for Spanish and Portuguese but failed to converge for Swedish. In both cases, the value of anchors in the REFINE step is clear, substantially improving the precision for all languages.

# SENTENCES	LANGUAGE MODEL	UAS / LAS		PERPLEXITY		ALIGN
		DEV	TEST	TRAIN	DEV	
28M	ELMo	72.3 / 62.8	72.5 / 61.3	22	44	85
10K	ELMo	52.9 / 38.3	50.1 / 33.1	4	4060	4
	ANCHORED ELMo	59.2 / 47.3	57.2 / 42.2	92	600	12

Table 4: Zero-shot, single-source results for the Spanish limited unlabeled data experiments. The parsing results are UAS/LAS scores, the perplexity is of the ELMo model, and the alignment scores are precision@5 on the held-out set, based on CSLS. All embeddings were aligned to English using supervised anchored alignment.

MODEL	LAS-F1
Rosa and Mareček (2018)	26.31
Smith et al. (2018a)	31.93
ALIGNED FASTTEXT	26.77
OURS	36.98

Table 5: Results for the Kazakh dataset from CoNLL 2018 Shared Task on Multilingual Parsing, compared to the two leading models w.r.t. this treebank.

Zero-Shot Parsing, Multiple Source Languages

Table 3 summarizes the results for our zero-shot, multi-source experiments on six languages from Google universal dependency treebank version 2.0.⁹ For each tested language, the parser was trained on all treebanks in the five other languages and English. We align each of the six languages to English. We compare our model to the performance of previous methods in the same setting (referred to as $L^t \cap L^s = \emptyset$ in Ammar et al. (2016)). The results show that our multilingual parser outperforms all previous parsers with a large margin of 6.8 LAS points. Even with an unsupervised alignment, our model consistently improves over previous models.

To make a fair comparison to previous models, we also use gold POS tags as inputs to our parser. However, for low-resource languages, we might not have access to such labels. Even without the use of POS tags at all, in five out of six languages the score is still higher than previous methods that do consider such annotations. An exception is the Portuguese language where it leads to a drop of 8.8 LAS points. While in the single language setting this good performance can be explained by the knowledge captured in the character level, contextual embeddings (Smith et al., 2018b; Belinkov

et al., 2017), the results suggest that this knowledge transfers across languages.

In order to assess the value of contextual embeddings, we also evaluate our model using non-contextual embeddings produced by FASTTEXT (Bojanowski et al., 2017). While these improve over previous works, our context-aware model outperforms them for all six languages in UAS score and for 5 out of 6 languages in LAS score, obtaining an average higher by 3 points. To further examine the impact of introducing context, we run our model with precomputed anchors (\bar{e}). Unlike FASTTEXT embeddings of size 300, these anchors share the same dimension with contextual embeddings but lack the contextual information. Indeed, the context-aware model is consistently better.

Few-Shot Parsing, Small Treebanks In this scenario, we assume a very small tree-bank for the tested language and no POS tags. We use the Kazakh tree-bank from CoNLL 2018 shared task (Zeman et al., 2018). The training set consists of only 38 trees and no development set is provided. Segmentation and tokenization are applied using UDPipe. Similar to Rosa and Mareček (2018); Smith et al. (2018a), we utilize the available training data in Turkish as it is a related language. To align contextual embeddings, we use a dictionary generated and provided by Rosa and Mareček (2018) and compute an alignment from Kazakh to Turkish. The dictionary was obtained using FastAlign (Dyer et al., 2013) on the Open-Subtitles2018 (Lison and Tiedemann, 2016) parallel sentences dataset from OPUS (Tiedemann, 2012).¹⁰

Table 5 summarizes the results, showing that our algorithm outperforms the best model from the shared task by 5.05 LAS points and improves by

⁹<https://github.com/ryanmcd/uni-dep-tb/>

¹⁰<https://github.com/CoNLL-UD-2018/CUNI-x-ling>

over 10 points over a FASTTEXT baseline.

Zero-Shot Parsing, Limited Unlabeled Data

To evaluate our anchored language model (Section 3.4), we simulate a low resource scenario by extracting only 10k random sentences out of the Spanish unlabeled data. We also extract 50k sentences for LM evaluation but perform all computations, such as anchor extraction, on the 10k training data. For a dictionary, we used the 5k training table from [Conneau et al. \(2018a\)](#).¹¹ Another table of size 1,500 was used to evaluate the alignment. In this scenario, we assume a single training language (English) and no usage of POS tags nor any labeled data for the tested language.

Table 4 shows the results. Reducing the amount of unlabeled data drastically decreases the precision by around 20 points. The regularization introduced in our anchored LM significantly improves the validation perplexity, leading to a gain of 7 UAS points and 9 LAS points.

7 Conclusion

We introduce a novel method for multilingual transfer that utilizes deep contextual embeddings of different languages, pretrained in an unsupervised fashion. At the core of our methods, we suggest to use anchors for tokens, reducing this problem to context-independent alignment. Our methods are compatible both for cases where a dictionary is present and absent, as well as for low-resource languages. The acquired alignment can be used to improve cross-lingual transfer learning, gaining from the contextual nature of the embeddings. We show that these methods lead to good word translation results, and improve significantly upon state-of-the-art zero-shot and few-shot cross-lingual dependency parsing models.

In addition, our analysis reveals interesting properties of the context-aware embeddings generated by the ELMo model. Those findings are another step towards understanding the nature of contextual word embeddings.

As our method is in its core task-independent, we conjecture that it can generalize to other tasks as well.

Acknowledgements

We thank the MIT NLP group and the reviewers for their helpful discussion and comments.

¹¹We filtered words with multiple translations to the most common one by Google Translate.

The first and third authors were supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract # FA8650-17-C-9116. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

This work was also supported in part by the US-Israel Binational Science Foundation (BSF, Grant No. 2012330), and by the Yandex Initiative in Machine Learning.

References

- Oliver Adams, Adam Makarucha, Graham Neubig, Steven Bird, and Trevor Cohn. 2017. [Cross-lingual word embeddings for low-resource language modeling](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 937–947. Association for Computational Linguistics.
- Hanan Aldarmaki and Mona Diab. 2019. [Context-aware crosslingual mapping](#). *arXiv preprint arXiv:1903.03243*.
- David Alvarez-Melis and Tommi Jaakkola. 2018. [Gromov-Wasserstein alignment of word embedding spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1881–1890. Association for Computational Linguistics.
- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah Smith. 2016. [Many languages, one parser](#). *Transactions of the Association for Computational Linguistics*, 4:431–444.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. [Learning bilingual word embeddings with \(almost\) no bilingual data](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462. Association for Computational Linguistics.
- Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2017. [Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1–10. Asian Federation of Natural Language Processing.

- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Wanxiang Che, Yijia Liu, Yuxuan Wang, Bo Zheng, and Ting Liu. 2018. [Towards better UD parsing: Deep contextualized word embeddings, ensemble, and treebank concatenation](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 55–64. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Christopher D. Manning, and Quoc Le. 2018. [Semi-supervised sequence modeling with cross-view training](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925. Association for Computational Linguistics.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018a. [Word translation without parallel data](#). In *International Conference on Learning Representations*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018b. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Georgiana Dinu and Marco Baroni. 2014. [Improving zero-shot learning by mitigating the hubness problem](#). *CoRR*, abs/1412.6568.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep Biaffine Attention for Neural Dependency Parsing](#). In *International Conference on Learning Representations*.
- Timothy Dozat, Peng Qi, and Christopher D. Manning. 2017. [Stanford’s graph-based neural dependency parser at the CoNLL 2017 Shared Task](#). In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 20–30. Association for Computational Linguistics.
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. [Cross-lingual transfer for unsupervised dependency parsing without parallel data](#). In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 113–122. Association for Computational Linguistics.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648. Association for Computational Linguistics.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6. Association for Computational Linguistics.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018a. [Learning word vectors for 157 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.
- Edouard Grave, Armand Joulin, and Quentin Berthet. 2018b. [Unsupervised alignment of embeddings with Wasserstein Procrustes](#). *CoRR*, abs/1805.11222.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2015. [Cross-lingual dependency parsing based on distributed representations](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1234–1244. Association for Computational Linguistics.
- 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 Thirtieth AAAI Conference on Artificial Intelligence*, pages 2734–2740.
- Yedid Hoshen and Lior Wolf. 2018. [Non-adversarial unsupervised word translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 469–478. Association for Computational Linguistics.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339. Association for Computational Linguistics.
- Chao Jiang, Hsiang-Fu Yu, Cho-Jui Hsieh, and Kai-Wei Chang. 2018. [Learning word embeddings for low-resource languages by PU learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1024–1034. Association for Computational Linguistics.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. [Loss in translation: Learning bilingual word mapping with a](#)

- retrieval criterion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2984. Association for Computational Linguistics.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. [Simple and accurate dependency parsing using bidirectional LSTM feature representations](#). *Transactions of the Association for Computational Linguistics*, 4:313–327.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. [Unsupervised machine translation using monolingual corpora only](#). In *International Conference on Learning Representations*.
- Angeliki Lazaridou, Georgiana Dinu, and Marco Baroni. 2015. [Hubness and pollution: Delving into cross-space mapping for zero-shot learning](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 270–280. Association for Computational Linguistics.
- Miryam de Lhoneux, Johannes Bjerva, Isabelle Augenstein, and Anders Søgaard. 2018. [Parameter sharing between dependency parsers for related languages](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4992–4997. Association for Computational Linguistics.
- Ying Lin, Shengqi Yang, Veselin Stoyanov, and Heng Ji. 2018. [A multi-lingual multi-task architecture for low-resource sequence labeling](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 799–809. Association for Computational Linguistics.
- Pierre Lison and Jrg Tiedemann. 2016. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. Linguistic knowledge and transferability of contextual representations. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. [Learned in translation: Contextualized word vectors](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6294–6305. Curran Associates, Inc.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. [Multi-source transfer of delexicalized dependency parsers](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 62–72.
- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. [Exploiting similarities among languages for machine translation](#). *CoRR*, abs/1309.4168.
- Tahira Naseem, Regina Barzilay, and Amir Globerson. 2012. [Selective sharing for multilingual dependency parsing](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 629–637. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. 2012. [A universal part-of-speech tagset](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*.
- Ye Qi, Devendra Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. [When and why are pre-trained word embeddings useful for neural machine translation?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 529–535. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#).
- Rudolf Rosa and David Mareček. 2018. [CUNI x-ling: Parsing under-resourced languages in CoNLL 2018 UD Shared Task](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 187–196. Association for Computational Linguistics.
- Aaron Smith, Bernd Bohnet, Miryam de Lhoneux, Joakim Nivre, Yan Shao, and Sara Stymne. 2018a. [82 treebanks, 34 models: Universal dependency parsing with multi-treebank models](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 113–123. Association for Computational Linguistics.
- Aaron Smith, Miryam de Lhoneux, Sara Stymne, and Joakim Nivre. 2018b. [An investigation of the interactions between pre-trained word embeddings, character models and pos tags in dependency parsing](#).

- In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2711–2720.
- Samuel L. Smith, David HP Turban, Steven Hamblin, and Nils Y Hammerla. 2017. [Offline bilingual word vectors, orthogonal transformations and the inverted softmax](#). In *International Conference on Learning Representations*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15:1929–1958.
- Milan Straka and Jana Straková. 2017. [Tokenizing, POS tagging, lemmatizing and parsing UD 2.0 with UDPipe](#). In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99, Vancouver, Canada. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*. European Language Resources Association (ELRA).
- Jörg Tiedemann. 2015. Cross-lingual dependency parsing with universal dependencies and predicted POS labels. In *Proceedings of the Third International Conference on Dependency Linguistics (Depling 2015)*, pages 340–349.
- Jörg Tiedemann. 2017. [Cross-lingual dependency parsing for closely related languages - helsinki’s submission to vardial 2017](#). In *Proceedings of the Fourth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial)*, pages 131–136. Association for Computational Linguistics.
- Dingquan Wang and Jason Eisner. 2016. [The galactic dependencies treebanks: Getting more data by synthesizing new languages](#). *Transactions of the Association for Computational Linguistics*, 4:491–505.
- Wenhui Wang, Baobao Chang, and Mairgup Mansur. 2018. [Improved dependency parsing using implicit word connections learned from unlabeled data](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2857–2863. Association for Computational Linguistics.
- Min Xiao and Yuhong Guo. 2014. [Distributed word representation learning for cross-lingual dependency parsing](#). In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 119–129. Association for Computational Linguistics.
- Min Xiao and Yuhong Guo. 2015. [Annotation projection-based representation learning for cross-lingual dependency parsing](#). In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 73–82. Association for Computational Linguistics.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. [Normalized word embedding and orthogonal transform for bilingual word translation](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011. Association for Computational Linguistics.
- Wei Yang, Wei Lu, and Vincent Zheng. 2017. [A simple regularization-based algorithm for learning cross-domain word embeddings](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2898–2904.
- Daniel Zeman, Jan Hajič, Martin Popel, Martin Potthast, Milan Straka, Filip Ginter, Joakim Nivre, and Slav Petrov. 2018. [CoNLL 2018 Shared Task: Multilingual parsing from raw text to universal dependencies](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 1–21, Brussels, Belgium. Association for Computational Linguistics.
- Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. [Adversarial training for unsupervised bilingual lexicon induction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1959–1970. Association for Computational Linguistics.
- Yuan Zhang and Regina Barzilay. 2015. [Hierarchical low-rank tensors for multilingual transfer parsing](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1857–1867. Association for Computational Linguistics.

A Alignment Results for All Layers

Table 6 manifests word-to-word translation results when supervised alignment is performed over different layer outputs from ELMo’s LSTM. Even though layer zero produces context independent representations, the anchors computed over the contextual representations achieved higher precision. We conjecture that this is due to the language model objective being applied to the output of the second layer. Hence, unlike token-based embeddings such as FASTTEXT that optimize them directly, the context-independent representations of ELMo are optimized to produce a good base for the contextual embeddings that are computed on top of them.

B Additional Parsing Results

In Table 7 we provide complementary results to those in zero-shot cross-lingual parsing.

LAYER	DE	ES	FR	IT	PT	SV	AVERAGE
0	54 / 71	65 / 80	66 / 80	61 / 78	55 / 71	41 / 61	57 / 74
1	62 / 78	73 / 85	74 / 86	69 / 82	61 / 74	49 / 68	65 / 79
2	59 / 75	68 / 82	70 / 83	66 / 79	56 / 72	48 / 67	61 / 76

Table 6: Per ELMo layer word translation to English precision @1 / @5 using CSLS (Conneau et al., 2018a) with a dictionary (supervised) for German (DE), Spanish (ES), French (FR), Italian (IT), Portuguese (PT) and Swedish (SV). Layer 0 representations are the result of the character-level word embeddings (which are context independent). Layer 1 and 2 alignments are based on anchors from the first and second LSTM layer output respectively.

MODEL	DE	ES	FR	IT	PT	SV	AVERAGE
Zhang and Barzilay (2015)	62.5	78.0	78.9	79.3	78.6	75.0	75.4
Guo et al. (2016)	65.0	79.0	77.7	78.5	81.9	78.3	76.7
ALIGNED FASTTEXT	69.2	83.4	84.6	84.3	86.0	80.6	81.4
ALIGNED \bar{e}	65.1	82.8	83.9	83.6	83.4	82.0	80.1
OURS	73.7	85.5	87.8	87.0	86.6	84.6	84.2
OURS, NO DICTIONARY	73.2	84.3	87.0	86.8	84.5	80.4	82.7
OURS, NO POS	69.7	84.8	85.3	85.3	79.7	81.7	81.1
OURS, NO DICTIONARY, NO POS	72.2	84.7	84.9	85.0	78.1	67.9	78.8

Table 7: Zero-shot cross lingual results compared to previous methods, measured in UAS. Aligned fastText and \bar{e} context-independent models are also presented as baselines. The bottom three rows are models that don’t use POS tags at all and/or use an unsupervised anchored alignment. Note that Ammar et al. (2016) did not publish UAS results.

C Hyperparameters

We now detail the hyperparameters used throughout our experiments. All alignment experiments were performed using the default hyperparameters of the MUSE framework (see their github repository). Table 8 depicts the values used in multilingual parsing experiments.

D Additional Alignment Example

We provide an additional example of a homonym. Figure 3 shows the contextual embeddings of the word “bank” in English and the words “banco” (a financial establishment) and “orilla” (shore) in Spanish. In this case, unlike the “bear” example (Figure 2), the embeddings do not form two obvious clusters in the reduced two dimensional space. A possible explanation is that here the two meanings have the same POS tag (Noun). Even so, as shown in Table 9, the alignment succeeds to place the embeddings of words from each context close to the matching translation.

The nearest-neighbors for the “bear” example are presented in Table 10.

HYPERPARAMETER	VALUE
BATCH SIZE (# SENTENCES)	32
INSTANCES PER EPOCH	32,000
EPOCHS (MAX.)	40
PATIENCE (EARLY STOPPING)	10
ENCODER TYPE	BI-LSTM
POS TAG EMBEDDING DIM.	50
LSTM, HIDDEN SIZE	200
LSTM, # LAYERS	3
DROPOUT RATE	0.33
ARC REPRESENTATION DIM.	500
TAG REPRESENTATION DIM.	100
ADAM	(DEFAULT)

Table 8: Hyper-parameters used in parsing experiments, shared across different settings.

<i>k</i>	“BANCO” ANCHOR
1	Unlike in primary succession , the species that dominate secondary succession , are usually present from the start of the process , often in the soil seed bank .
2	Canto XLV is a litany against Usura or usury , which Pound later defined as a charge on credit regardless of potential or actual production and the creation of wealth ex nihilo by a bank to the benefit of its shareholders .
3	This prompted some investigation , led by Sir Benjamin Hall , which quickly turned up the fact that O’ Connor was registered as the owner of all the estates , and of the associated bank .
4	The commercial NaS battery bank offers : (Japanese) .
5	Both team leaders are given a mystery word , which along with their team - mates use gigantic foam blocks and place them on the clue bank (similar to Boggle) with only giving a clue to the word ...
<i>k</i>	“ORILLA” ANCHOR
1	The combined Protestant forces , now numbering 25,000 strong , positioned themselves on the western bank of the Rhine River .
2	The Romans had a small advance guard of auxiliaries and cavalry on the opposite bank of the river .
3	Between Vaugirard and the river Seine he had a considerable force of cavalry , the front of which was flanked by a battery advantageously posted near Auteuil on the right bank of the river .
4	Mallus therefore stood on the eastern bank of the river .
5	The Argentine squadron spent the night of February 7 anchored between Juncal Island and the west bank of the river .
6	After marching north from Tewkesbury , Sir William Waller tried to contain the cavalry forces of Maurice on the western bank of the Severn , cutting this substantial force off from the rest of the Royalist army .

Table 9: Nearest-neighbors (after alignment) of the Spanish anchors “banco” (a financial establishment) and “orilla” (shore) from the contextual embeddings of the word “bank” in English. The full sentence is presented for context.




Figure 3: Contextual embeddings for the English word “bank” and its two possible translations in Spanish — “banco” (a financial establishment) in red and “orilla” (shore) in blue. The figure shows a two dimensional PCA for the aligned space of the two languages. The symbols are the anchors and the dots are the contextualized embeddings. (best viewed in color)

<i>k</i>	“TENER” ANCHOR
1	It may be difficult for the patient to bear the odour of the smoke at first , but once he gets used to such a smell , it does not really matter .
2	No matter what the better class of slave owners might do , they had to bear the stigma of cruelty with the worst of tyrants ...
3	Every new car will bear the lan name instead of Van Diemen , so the highly successful marque will gradually disappear .
4	had a sufficient economic stake to bear the litigation burden necessary to maintain a private suit for recovery under section 4 .
5	In this example , consumers bear the entire burden of the tax ; the tax incidence falls on consumers .
<i>k</i>	“OSO” ANCHOR
1	In 2010 , the government of the NWT decided to update its version of the polar bear - shaped plate .
2	Salad Fingers appears to be masochistic , as he can be seen taking pleasure from impaling his finger on a nail , rubbing stinging nettles on himself or stepping onto a bear trap .
3	The old bear - hunter , on being toasted , made a speech to the Texians , replete with his usual dry humor .
4	Balto arrives , distracts the bear , saves Aleu , they both escape and the bear disappears .
5	Defeated , the polar bear shrinks and transforms into a plush toy .

Table 10: Nearest-neighbors (after alignment) of the Spanish anchors “tener” (carry, verb) and “oso” (animal, noun) from the contextual embeddings of the word “bear” in English. The full sentence is presented for context.