

Towards a Broad Coverage Named Entity Resource: A Data-Efficient Approach for Many Diverse Languages

Silvia Severini, Ayyoob Imani, Philipp Dufter*, Hinrich Schütze

Center for Information and Language Processing

LMU Munich, Germany

silvia@cis.uni-muenchen.de

Abstract

Parallel corpora are ideal for extracting a *multilingual named entity (MNE) resource*, i.e., a dataset of names translated into multiple languages. Prior work on extracting MNE datasets from parallel corpora required resources such as large monolingual corpora or word aligners that are unavailable or perform poorly for under-resourced languages. We present CLC-BN, a new method for creating an MNE resource, and apply it to the Parallel Bible Corpus, a corpus of more than 1000 languages. CLC-BN learns a neural transliteration model from parallel-corpus statistics, without requiring any other bilingual resources, word aligners, or seed data. Experimental results show that CLC-BN clearly outperforms prior work. We release an MNE resource for 1340 languages and demonstrate its effectiveness in two downstream tasks: knowledge graph augmentation and bilingual lexicon induction.

Keywords: Low-resource, Multilinguality, Named Entities, Transliteration

1. Introduction

Of the thousands of languages in the world, a very small portion is covered by language technologies (Joshi et al., 2020). Bird (2020) suggests a number of approaches to develop such technologies for low-resource languages. In this paper, our goal is to create a *multilingual named entity (MNE) resource* – by which we mean a dataset of names translated into multiple languages – for a large number of low-resource languages, in total more than a thousand. Named entities (NEs) are crucial for many language technologies and NLP applications, including text comprehension, question answering, information retrieval and relation extraction. In this paper, we demonstrate the effectiveness of our MNE resource in two downstream tasks: knowledge graph augmentation and bilingual lexicon induction.

We extract our MNE resource from the Parallel Bible Corpus (PBC) (Mayer and Cysouw, 2014), a multiparallel corpus that covers more than 1300 languages. (Note however that we do not use Bible-specific features; therefore, our work is in principle applicable to any parallel corpus.) For some languages, PBC is the only available text (Wu et al., 2018). Multiparallel corpora contain sentence-level parallel text in more than two languages. Apart from PBC, JW300 (Agić and Vulić, 2019) and Tatoeba¹ are two other examples of such corpora. While the amount of data per language provided by highly multiparallel corpora is usually small, this type of corpus plays an important part in compiling resources for low-resource languages.

Creating a named entity resource is comparatively easy if sufficiently high-quality resources are available for




Figure 1: Two NEs from our resource, each with a sample of translations in six different languages

a language. Such resources include named entity recognizers (Yadav and Bethard, 2018; Li et al., 2020); large monolingual corpora, which can be used to learn high-quality word embeddings or high-quality contextualized embeddings; parallel corpora that consist of large corpora (millions of words) per language (Lample et al., 2016; Ma and Hovy, 2016; Dasigi and Diab, 2011); or high-quality annotated data, e.g., training sets for named entity recognition (Wang and Manning, 2014; Wu et al., 2021; Wu et al., 2020) or implicit high-quality annotations like hyperlinks in Wikipedia (Tsai et al., 2016). Recent work (Wu et al., 2021; Li et al., 2021) with multilingual pretrained language models (PLMs) like BERT and XML-R for named entity recognition is promising, but also relies on moderately large monolingual corpora (e.g., a Wikipedia of decent size) to learn good quality contextualized representations. However, these monolingual corpora exist only for about 100 or so languages. For instance, Zulu is not included but we cover it in our experiments.

In this work, our goal is to cover the large number of languages for which these resources do not exist: no named entity recognizers, no large monolingual (or parallel) corpora, no annotated data (not even implicitly

*Now at Apple.

¹<https://tatoeba.org>

annotated) and no pretrained language models (due to the lack of large monolingual corpora).

Many low-resource languages are covered in the PBC which gives us a chance to create resources for languages that currently do not have any – perhaps apart from an entry in the World Atlas of Language Structures (Dryer and Haspelmath, 2013) that is too abstract for most purposes in computational linguistics.

Since PBC is a parallel corpus, the question of why we do not use word alignment naturally arises. However, our experiments with word alignment on PBC were not successful for named entities. The reason is that word alignment performance deteriorates when parallel text is scarce (Och and Ney, 2003), especially for named entities as most are rare words. Our approach therefore does not depend on a word aligner and works well even when only a small parallel corpus is available. We directly compare with prior work that relies on word alignment. Based on this motivation, we introduce CLC-BN (*Character Level Correspondence Bootstrapping and Neural transliteration*), a method for extracting a multilingual named entity resource from a parallel corpus, including in low-resource settings in which the available text per language in the corpus is small. CLC-BN learns a neural transliteration model from parallel-corpus statistics, without requiring any other bilingual resources, word aligners or seed data. In the first step, the method identifies NE correspondences in the parallel text. It then learn a neural transliteration model from these (noisy) NE correspondences. Finally, we use the learned model to identify high-confidence NE pairs in the parallel text. The first step (identifying NE correspondences) works at the character-ngram level, hence it is applicable to languages for which a tokenizer is not available, as opposed to word alignment based approaches. We will show that our method performs well for untokenized Japanese text.

In summary, our contributions are:

1. We present CLC-BN, a method that first identifies named entity correspondences in a parallel corpus and then learns a neural transliteration model from them.
2. We annotate a set of NEs to evaluate CLC-BN’s performance on 13 languages through crowdsourcing and show a clear performance increase in comparison to prior work. We release the gold annotated sets as a resource for future work.²
3. Using CLC-BN, we create and release a named entity resource containing 674,493 names across 1340 languages, 503 names per language on average.²
4. For many languages, ours is the first published resource. We believe that it can be useful for future work in computational linguistics on the more than

1000 languages covered. We show experimentally that this is the case for knowledge graph augmentation and bilingual lexicon induction.

2. Related work

2.1. Word alignment

A multilingual named entity resource can be extracted from a parallel corpus via word alignment. Word alignment has been widely studied. Statistical word alignment models were introduced by Brown et al. (1993). More recently Giza++ (Och and Ney, 2000) and Eflomal (Östling et al., 2016) were released followed by neural network extensions (Ngo-Ho and Yvon, 2019). Other approaches use learned representations for creating alignments (Jalili Sabet et al., 2020). In concurrent work, Imani et al. (2021) have shown that better word alignment results can be achieved by exploiting multi-parallel corpora. Previous work on named entity alignment and recognition uses combinations of alignment tools and postprocessing techniques. Dasigi and Diab (2011) use Giza++ for alignment and applied statistical machine translation (Koehn et al., 2007) and language-specific rules for improving transliteration. (Wu et al., 2018) use the Berkeley aligner (Liang et al., 2006) to word-align language pairs in the English Bible and further improve them with machine translation. In this paper, we do not use word aligners because of their low quality for named entities in small parallel corpora. We will directly compare with the word-alignment-based method of (Wu et al., 2018).

Recent approaches rely on parallel corpora and multilingual pre-trained models. Wu et al. (2021) construct a pseudo training set by performing translation and use multilingual BERT (Devlin et al., 2019a) to generate language independent features for training NER models. Li et al. (2021) use XLM-R (Conneau and Lample, 2019) to build an entity alignment model that projects English named entities into the parallel target language. While these approaches are promising, they are limited to the language set the models have been trained on (≈ 100). In contrast we apply CLC-BN to the more than one thousand languages in the Parallel Bible Corpus.

2.2. Transliteration

Prabhakar and Pal (2018) provide a comprehensive survey on transliteration. Recently, the task has been addressed with sequence-to-sequence models and transformers. Wu and Yarowsky (2018) perform experiments with these models on their Bible-based translation matrix dataset (Wu et al., 2018) and show that the task is challenging in the low-resource scenario. One of the causes is overfitting of the training set due to its reduced size. Our CLC-BN method uses a transliteration model and addresses this problem by augmenting the training set with monolingual target data (English) and introducing a monotonic bias.

²http://cistern.cis.lmu.de/ne_bible/

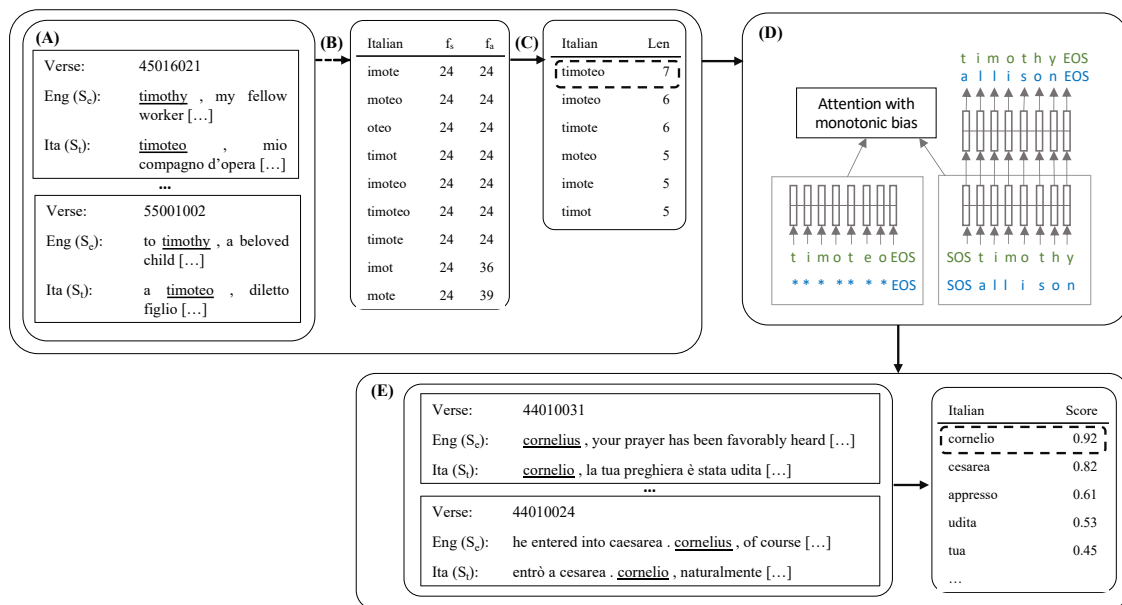


Figure 2: Data flow in CLC-BN. Example showing extraction of Italian NE training candidates for English “timothy” and identification of an Italian NE that matches English “cornelius”. The input is the parallel corpus (A). CLC-B extracts from the parallel corpus ngrams that are candidate transliterations for “timothy” (B). These candidates are then filtered (C). (D): The architecture of the neural transliteration model. Green input-output pairs: Italian-English training data taken from the output of CLC-B. Blue input-output pairs: monolingual English training data. (E): We use the trained neural model to score candidates taken from the Italian parallel verses in which “cornelius” appears and keep the best scoring word.

2.3. Named entity resources

(Benites et al., 2020) introduce Translit, a transliteration resource created by combining and unifying public corpora. However, this dataset only covers 180 languages. BabelNet (Navigli and Ponzetto, 2012) is a multilingual encyclopedic dictionary that integrates WordNet, Wikipedia, GeoNames, inter alia. BabelNet is more comprehensive than other resources, but its NE coverage is still poor for many languages (e.g., for Inuktitut). We show in this paper that we can extend the coverage of BabelNet with our method. The Translation Matrix of (Wu et al., 2018) covers 591 languages. Their approach is based on word alignment. We show that our approach yields higher quality.

2.3.1. Named Entity Recognition resources

Named Entity Recognition (NER) systems usually require annotated data to achieve high accuracy. Our NE resource can be exploited to bootstrap such NER models for many different languages. (Al-Rfou et al., 2015) automatically extract named entities from Wikipedia link structure and Freebase attributes and create PolyglotNER for 40 languages. (Pan et al., 2017) introduce WikiAnn, a resource for 282 Wikipedia languages that supports name tagging and entity linking. Our resource covers more than 1300 languages and CLC-BN does not rely on external sources other than the PBC.

2.4. Annotation projection

(Ehrmann et al., 2011) project annotations from English to five languages using a phrase-based statistical ma-

chine translation system and different methods: string matching, consonant signature matching and edit distance similarity. Ni et al. (2017) propose two methods for NER projection using heuristics, alignment information, and mapped word embeddings. Wang et al. (2018) describe a method for cross-lingual knowledge graph alignment of pre-aligned entities based on their distance in the learned embedding space. We project English NEs to the target languages exploiting character-level correspondence and a neural transliteration model without requiring any word alignment information or seed data.

2.5. Monotonicity

The performance of sequence-to-sequence models on some tasks can be improved by imposing an inductive bias of monotonicity (i.e., no character can be aligned to one that precedes a previously aligned character). Previous studies implement and analyze the effect of such a monotonic bias. Wu and Cotterell (2019) show that enforcing strict monotonicity and learning a latent alignment jointly while learning to transduce leads to improved performance for morphological inflection, transliteration, and grapheme-to-phoneme conversion. Rios et al. (2021) develop a general method for incorporating monotonicity into attention for seq2seq and Transformer models, agnostic of the task and model architectures. Similar to this prior work, we impose a monotonic bias on our neural transliteration model.

3. Method

We now describe CLC-BN.³ Figure 2 shows architecture and data flow. For ease of development and evaluation, we also use the Uroman romanizer (Hernjakob et al., 2018). It converts scripts into Latin characters. But CLC-BN can be applied equally well without romanization. CLC-BN consists of two steps. First we extract character-level correspondences (CLC-B). Then we train a neural transliteration model to obtain the final set of named entities.

3.1. Character-Level Correspondence Bootstrapping (CLC-B)

We use cooccurrence statistics at the character level between English NEs and target language NEs to create a training set for the neural transliteration model. We use (Wu et al., 2018)’s list of English Bible NEs. NEs with frequency 1 are not considered in CLC-B because the FILTER step (#3 below) is likely to produce false positives (accidentally correlated ngrams) for them; but they are considered in §3.2.

CLC-B is designed based on the following simple correspondence assumption: if an English NE occurs in a verse, the corresponding target NE occurs in the parallel target verse and vice versa. This also implies that if N and M are the the number of verses in which the NE and its translation occur, then $N \approx M$. We do not require $N = M$ because we relax the correspondence assumption due to errors in the parallel corpus and due to the use of pronouns (including null pronouns, i.e., the pronoun is only present implicitly), which differs across languages.

We now describe our Character-Level Correspondence Bootstrapping (CLC-B) method, for the example of an English NE w . Algorithm 1 shows the pseudocode. Let f_a be the total frequency of an ngram in the target language and f_s its frequency in the subset of verses that contains w in English.

1. **EXTRACT.** (Line 4) Extract the parallel subcorpus that contains w from the parallel corpus. It consists of the English part S_e and the target language part S_t .
2. **GET_NGRAMS.** (Lines 5–13) For all character n -grams⁴ ($3 < n < 20$) in S_t , determine f_s , the number of occurrences in S_t . Discard ngrams with $f_a > 50$ – this removes a small number of frequent NEs like Jesus, but avoids false positive matches with frequent ngrams. The resulting set of target ngrams is G_t .
3. **FILTER.** (Line 14) Filter G_t as follows. (a) Determine the ngram(s) with the highest f_s . Remove all other ngrams. (b) Determine the ngram(s) with the minimum absolute difference between f_a and f_s .

³Reproducibility details in §A.

⁴We discard ngrams containing digits, punctuation and spaces.

Algorithm 1 Pseudocode for the CLC-B method. Given a parallel corpus of English (E) and a target language (T), we identify, for each English NE, its target match. See §3.1 for details and for the EXTRACT and FILTER methods.

```
1: procedure CLC-B(corpus  $E$ , corpus  $T$ , list  
   English_NEs)  
2:    $pairs \leftarrow list()$   
3:   for  $w \in$  English_NEs do  
4:      $S_e, S_t \leftarrow extract(w, S, T)$   $\triangleright$  (1) EXTRACT  
5:      $G_t \leftarrow list()$   
6:      $ngram\_list \leftarrow get\_ngram\_list(S_t)$   
7:      $frequency\_list \leftarrow get\_frequent\_ngrams(S_t)$   
8:     for  $[ngram, count] \in ngram\_list$  do  
9:       if  $ngram \in frequency\_list$  or  $count == 1$  then  
10:        continue  
11:      end if  
12:       $G_t.append([ngram, count])$   
13:    end for  
14:     $pairs.append(filter(G_t))$   $\triangleright$  (3) FILTER  
15:  end for  
16:  return pairs  
17: end procedure
```

Remove all other ngrams. Intuitively, most NEs in a particular domain are unique – so they should contain ngrams that only occur in the NE and not in other words. (c) Return the ngrams with the smallest length difference to w . This eliminates candidates that are much longer or shorter than w .

3.2. Neural transliteration

CLC-B returns a noisy set of NE pairs, especially when only a small number of parallel verses is available for a language (we refer to this as the *lowest-resource* setting below). We build a neural sequence-to-sequence model (Sutskever et al., 2014) to refine it and to mine additional pairs. We use a single-layer bidirectional Gated Recurrent Unit (GRU) (Cho et al., 2014) encoder and a single-layer GRU decoder with attention (Luong et al., 2015). The sequences are processed at the character-level, with separate input and output vocabularies. Target language NEs are the input, English NEs the output; we use input/output when referring to the neural model (not source/target) because “target” in this paper refers to the target language that English is paired with.

To make best use of the limited training data in our experimental setup, we use augmentation and impose a monotonicity bias as described below. To avoid overfitting, we augment the training set with English NEs. We label the English Wikipedia dump⁵ with the Flair Part-of-Speech tagger (Akbik et al., 2019), and select all NEs. We add, for each English NE mined from Wikipedia, one pair of the form (empty input NE, English output NE) to the training set. We use empty input NEs to prevent the learning of the identity function while helping the decoder to learn the structure of English words. To

⁵<https://dumps.wikimedia.org/> (01.04.2020)

	Lang	ISO	# verses	# parallel
low-resource languages	Arabic	Arb	31173	31062
	Finnish	Fin	31167	31061
	Greek	Ell	31183	31062
	Russian	Rus	31173	31062
	Spanish	Spa	31167	31062
	Swedish	Swe	31167	31062
	Zulu	Zul	31167	31062
lowest-resource languages	Hebrew	Heb	7952	7917
	Hindi	Hin	7952	7917
	Kannada	Kan	7952	7917
	Korean	Kor	7913	7869
	Georgian	Kat	4904	4844
	Tamil	Tam	7942	7917

Table 1: Number of verses in PBC and number of verses that are parallel with our English edition for the languages in our experiments. The English edition has 31,133 verses.

prevent generation of output independent of the input, we ensure equal proportions of original and augmented data by oversampling the former. Because transliterations are (with few exceptions) monotonic, we impose a monotonicity bias: we mask the attention matrix, so that the model cannot see anything to the left of the position previously attended to.

Given an English NE w and the verses S_e in which it appears, target candidates are all words in S_t , the verses parallel to S_e . Once the model is trained, we choose the best scoring candidate as w 's transliteration where the score is the average log likelihood of the output characters (Severini et al., 2020).

We use a slightly different scoring step for non-tokenized languages (e.g., Japanese) because separated words in S_t are not available: given an English NE w , the target candidates are all ngrams that CLC-B has extracted for w in step 3b.

4. Evaluation and Analysis

We apply CLC-BN to the Parallel Bible Corpus (PBC) (Mayer and Cysouw, 2014) for evaluation and for creating our NE resource.⁶ We evaluate on a subset of 13 languages that includes different scripts, resource availabilities and language families: Arabic, Greek, Finnish, Hebrew, Hindi, Kannada, Korean, Georgian, Russian, Spanish, Swedish, Tamil, and Zulu. These languages are also covered by the baselines and are therefore suitable for comparison. We view them as a representative subset for evaluating our method's performance. Note, however, that our NE resource covers all 1340 PBC languages: our approach is applicable to all languages since it does not use language-specific features and pre-processing steps.

PBC contains 1340 languages, most of which are low-resource. It is divided into subfiles, each containing Bible text from one language. Some languages that cover the Hebrew Bible and the New Testament com-

pletely contain about 30,000 verses. Other languages contain fewer than 8000 verses. We divide the languages into two categories: **lowest-resource**, fewer than 8000 verses; and **low-resource**, between 8000 and 32,000 verses.⁷ Table 1 gives the number of verses for the editions we use. We evaluate our resource on human annotated data and on silver data with respect to the baselines and provide analysis.

4.1. Human evaluation

We annotated 60 NEs per language using Toloka,⁸ a crowd-sourcing platform. Annotators had to pass an English test and successfully complete a training task to gain access to the annotation pool. Their performance was constantly checked using covert control questions. Each question contained the English NE and up to five possible options: one for each of the three baselines, one for CLC-B and one for CLC-BN. Each option consists of the word in the target script together with its romanized version in parentheses. Annotators had to mark all correct options that can be paired to the English NE, or none if no option is correct. Each question was annotated by exactly three annotators.

We calculate annotator agreement using Cohen's Kappa (Cohen, 1960), which measures agreement above chance. Similar to the setup of (Wu et al., 2018), we do not require that the annotators know the target languages. However, their average pairwise agreement is 0.73, "substantial agreement" according to Cohen's Kappa (Landis and Koch, 1977), indicating that they can find the correct corresponding target named entity even if they do not know the target language. To create the final gold set, we adopt a majority voting strategy and keep named entities that at least two annotators agreed on, resulting in at least 58 named entities per language. We evaluate CLC-BN and the baselines on this gold set.⁹ The results can be found in Table 2, column "Hum". CLC-BN outperforms the baseline (Wu et al., 2018) for all languages (average difference of 7.9), with substantial improvements for the lowest-resource languages (difference of 21.1). The biggest improvements are for Hindi and Kannada (more than 30).

4.2. Silver evaluation

The gold dataset is used as the main evaluation of the resource. However, we additionally create a silver dataset to evaluate based on a larger set of hundreds of NEs. We create the silver set by translating each English NE to all target languages supported by the Google translation API¹⁰ and comparing them with the NEs extracted

⁷"low-resource" is to be interpreted as referring to the setting in our experiments. For example, many resources are available for Russian, but in our setting we only use the Russian text that is available in PBC to evaluate how well our method works in a low-resource setting.

⁸<https://toloka.yandex.com/>

⁹We release the gold dataset to facilitate future research.

¹⁰<https://cloud.google.com/translate>

⁶Reproducibility details in §A

	Arb	Ell	Fin	Spa	Swe	Rus	Zul	AVG
	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum
(Wu et al., 2018)	67.9 70.0	47.2 80.0	<u>89.0</u> 90.0	87.6 91.7	88.8 88.3	60.6 72.9	61.9 84.8	65.9 82.5
Östling et al. (2016)	69.8 61.7	53.4 88.3	77.7 76.7	83.9 86.7	81.2 85.0	64.8 83.1	52.9 86.4	60.9 81.1
Sabet et al. (2020)	18.1 20.0	23.5 40.0	49.8 60.0	35.6 45.0	41.6 50.0	39.6 45.8	18.3 25.4	29.6 40.9
CLC-B	53.8 56.7	32.2 45.0	59.9 50.0	48.0 48.3	52.0 48.3	46.5 57.6	55.1 74.6	46.6 54.4
CLC-BN	<u>70.6</u> <u>81.7</u>	<u>54.7</u> <u>91.7</u>	86.5 <u>93.3</u>	<u>89.6</u> <u>96.7</u>	<u>89.9</u> <u>91.7</u>	<u>70.2</u> <u>84.8</u>	<u>68.8</u> <u>93.2</u>	<u>71.9</u> <u>90.4</u>
	Heb	Hin	Kan	Kat	Kor	Tam	AVG	
	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum	Dist Hum
(Wu et al., 2018)	53.4 62.5	64.1* 76.3*	41.5 61.7	64.3 70.0	30.5 54.2	47.1* 66.1*	50.2 65.1	
Östling et al. (2016)	<u>65.5</u> <u>83.9</u>	57.7* 69.5*	23.1 38.3	64.0 68.3	16.6 33.9	20.5* 35.6*	41.2 57.5	
Sabet et al. (2020)	27.8 23.2	41.6* 47.5*	26.5 46.7	28.0 20.0	23.5 40.0	30.4* 47.5*	29.6 37.5	
CLC-B	37.2 51.8	43.1* 39.0*	30.7 48.3	41.7 45.0	23.8 37.3	38.6* 47.5*	35.9 44.8	
CLC-BN	62.6 71.4	<u>78.6*</u> <u>94.9*</u>	45.9 <u>93.3</u>	<u>70.8</u> <u>88.3</u>	<u>34.2</u> <u>78.0</u>	<u>59.5*</u> <u>91.5*</u>	<u>58.6</u> <u>86.2</u>	

Table 2: Precision of NE correspondence identification for low-resource (top: Hebrew Bible and New Testament) and lowest-resource (bottom: New Testament only) languages. We compare Translation Matrix (Wu et al., 2018), Eflomal (Östling et al., 2016), SimAlign (Sabet et al., 2020), CLC-B and CLC-BN. Comparisons are with silver data+Jaro distance (Dist) and with gold human annotated data (Hum). *: evaluation on romanization for fair comparison with baselines.

by CLC-BN using the Jaro distance (Jaro, 1989). The distance takes into account the number and order of characters shared by two strings; e.g., the NE “salome” has a distance of 0.05 from “salom” and 0.11 from “calom”. Jaro is frequently used for entity matching and is well-suited for short strings (Cohen et al., 2003). We use a threshold of 0.3 for the Jaro distance, chosen to be strict enough to evaluate the NEs and to take into account noise in the pairs produced by Google Translate. For example, the silver translation of “jannes” in Greek is γιάννης (giánnēs) while our data contains ιαννής (iannís), which is also correct; their distance is 0.26. Another example is the name “mitylene” that the silver data translates to μυτυλένιο (mytylénio) and has a distance of 0.27 to our translation μυτιλήνη (mytilíni). By design, the silver data provides only a single translation for each English NE. However, multiple translations are often correct, due to the variability of morphology, transliteration, naming conventions and dialects (Prabhakar and Pal, 2018). For example, the English NE “Paul” can be aligned to Russian “Pavel” and “Pavla”. For this reason, our results on the silver standard must be interpreted as lower bounds.

Arabic and Hebrew are standardly written without short vowels. This is also the case for the silver data. However, some PBC editions are written with short vowels, so we postprocess predictions by removing short vowel diacritics.

Table 2 shows results for the 13 languages. The ranking of baselines and methods is similar to the one obtained with the gold human evaluation with CLC-BN being always the best, except for Finnish. Improvements for lowest-resource languages (lower part of the table) are large, up to 48% difference on average. CLC-BN outperforms (Wu et al., 2018) for 12 of the 13 languages.¹¹

¹¹The exception is Finnish, which is probably due to the fact that machine translation (which was used for (Wu et al., 2018)) performs well for high-resource languages. Note however, that CLC-BN performs best for Finnish in the (more reliable) human (“Hum”) evaluation.

4.3. Word alignment comparison

NE correspondences can also be obtained using a word aligner. We compare our results with pairs obtained using Eflomal (Östling et al., 2016), a statistical word aligner, and SimAlign (Sabet et al., 2020), a high-quality word aligner that leverages multilingual word embeddings. Table 2 shows precision for silver and gold data. CLC-BN outperforms Eflomal (with the exception of Hebrew) and SimAlign for all 13 languages. We attribute this to the fact that NEs are hard to word-align because most of them are infrequent, resulting in alignment errors due to sparseness. CLC-BN could be integrated into word alignment pipelines to boost word aligner performance for NEs (Sajjad et al., 2011; Semmar and Saadane, 2013).

CLC-B works at the character level and is applicable to non-tokenized languages while aligners are not. Japanese is non-tokenized, so we evaluate it (only for CLC-B and CLC-BN since the other methods were not run on Japanese). We evaluate the 979 pairs of CLC-BN with the silver data and obtain a precision of 63.2%. We also use Toloka for the gold evaluation of 60 random pairs and obtain a precision of 60%. However, in this case each question has at most two options (CLC-B and CLC-BN – in contrast to five as for 4.1), which can hinder the annotators’ judgments having less comparison terms. For this reason, we also asked three experts to evaluate the 60 pairs and obtained a precision of 85%.

4.4. Impact of corpus size

Table 2 shows that precision for lowest-resource languages (less than 8000 verses, bottom) is worse than those for low-resource languages (about 30,000 verses, top), with an average difference of 13.3% for silver data, and 4.2% for gold data. The small gap on gold data, highlights that our method is appropriate also for the lowest-resource setting. Table 3 shows some examples of aligned pairs according to CLC-BN. We see that errors arise as the frequency of NEs in the English corpus diminishes. For example, the Kannada alignment for

#	English	Arabic	Finnish	Greek	Hebrew	Kannada	Russian	Tamil
28	elijah	alalihaau	eliaa	elia	veaeliyahu	eliiyanaagali	elisei	eliyaavaa
12	titus	tiytusa	titus	titos	titos	titanannu	titu	tiittuvin
8	elizabeth	aaliysaabaata	elisabet	elisabet	elisheva	elisabeet	elizaveta	elicapet
3	miletus	milyiyusa	miletokseen	miletu	lemilitos	mileetakke	mileta	mileettu
2	rufus	ruwfusa	rufuksen	roufo	vishelom	uphaniguu	rufa	ruupuvukkum
2	hermes	wahirmisa	hermeeksi	epairne	heremes	meeyaniguu	germes	ermee

Table 3: Examples of named entity alignments (romanized). “#” column shows the number of verses in which the English word appears.

Lang	Eng	Freq	CLC-B	CLC-BN
Arb	anah	10	الشيخة (alshaykha)	لنا (ana)
Rus	joanna	2	мария (mariya)	иоанна (joanna)
Fin	perez	2	hesroni	peresin
Kan	cainan	2	ನಾನಾ (naanaa)	ಕಯಿನಾನಾ (kayinaanana)
Tam	azor	2	எலியூதுக்குத் (eliyutukku)	ஆசோர் (aacoor)

Table 4: Examples of improvement due to neural transliteration. CLC-B: incorrect prediction of CLC-B. CLC-BN: correct prediction obtained with neural transliteration.

“rufus” and Greek and Kannada alignments for “hermes” are incorrect. Both words are short, indicating another source of errors: short words provide less of a signal for the neural transliteration model than long words do.

4.5. Impact of neural transliteration

Table 2 shows precision for CLC-B and CLC-BN. All languages benefit from neural transliteration with an average improvement of 30.9 percentage points. One of the reasons is that CLC-B was designed to discard English NEs that appear only once in the corpus. Table 4 shows examples where neural transliteration corrects an error made by CLC-B. Most of these cases have low frequency. This is not surprising as the risk of false positives increases as the frequency decreases because the heuristics used in CLC-B (§3.1) are less reliable for low-frequency NEs.

4.6. Error analysis

In our manual error analysis, we found two main types of errors.

(1) The neural model generally learns well how to transliterate the beginning of a word, but error rates are higher word-internally. For example, the NE “balak” is wrongly paired to “pileeyaam” instead of “paalaak” and “menna” is paired to “meleyaa” instead of “meyinaan” in Tamil. The neural model has to learn two aspects of transliteration: transliteration proper (i.e., character correspondences) and alignment. This type of error indicates that alignment performance should be improved. In future work, we plan to explore neural architectures that more explicitly model the problem as alignment.

(2) For some low-resource languages, the output of CLC-B has a high level of noise, so the neural model fails to

learn some character correspondences. In some cases, the output of the neural model is unrelated to the input. This type of error indicates that the CLC-B method should be improved further. As shown in Tables 3 and 4, low-frequency words contain more errors. In future work, we plan to adopt an iterative strategy that considers gradually more and more named entities, starting with the most confident ones.

5. Use cases

5.1. Transliteration

A straightforward application of our named entity resource, as described by (Wu et al., 2018), is to create transliteration models. They showed that a character-based Moses SMT system trained over a dataset of named entities extracted from the Bible (whose performance is lower than our method’s, based on Table 2) performs better than a Unicode baseline. We now present two additional applications of our named entity resource: extending existing multilingual dictionaries and cross-lingual mapping of word embeddings.

5.2. Extending existing multilingual resources

BabelNet¹² (Navigli and Ponzetto, 2012) is a multilingual encyclopedic dictionary. It was created by integrating more than 35 WordNets, covering 500 languages, and has about 20 million entries.

We want to show that one can use our resource to enrich BabelNet further. Since CLC-BN covers many more languages than BabelNet, we can simply extend BabelNet by adding more languages like Burarra, North Junín Quechua, and Mian to it. Regarding the languages that BabelNet already supports, we check whether we can add more entries exploiting our resource. To this end, for each word pair (English:target-language) in CLC-BN, we check whether a translation of the English word exists in BabelNet in the target language. Results are depicted in Table 5. On average, 27% (i.e., 206 words) of the English words have no correspondence in the target language. These are mostly rare words that are difficult to translate without accessing a resource as rich as PBC. From a manual investigation, we find that our resource could also help to improve the quality of BabelNet; some translations of the latter are completely incorrect or wrongly written with Latin characters. Examples for Greek are hamor/εμμωρ (emmor), which

¹²<https://babelnet.org/>

Lang.	CLC-BN	Babel	New NEs	New NEs %
Arb	977	683	294	30.1
Fin	979	647	332	33.9
Ell	979	658	321	32.8
Rus	485	449	36	7.4
Spa	979	784	195	19.9
Swe	979	684	295	30.1
Zul	979	471	508	51.9
Heb	467	413	54	11.6
Hin	467	334	133	28.5
Kan	467	299	168	36.0
Kor	467	386	81	17.3
Kat	368	271	97	26.4
Tam	433	318	115	26.6
Jpn	979	715	264	27.0
Zho	979	698	281	28.7
Tha	467	337	130	27.8
AVG.	715	509	206	27.2

Table 5: Extension of BabelNet with named entities based on our resource. Example (first line, “Arb”). CLC-BN returns 977 English-Arabic NE pairs. BabelNet contains Arabic translations for 683 of these English NEs, but 294 (30.1%) lack an Arabic translation. Thus we add 294 English-Arabic NE pairs that were not covered by BabelNet.

BabelNet translates as Δείνα (Deina), and ethan/εθόλν, incorrectly transliterated with Latin characters.

5.3. Cross-lingual mapping of word embeddings

An effective method for creating bilingual word embeddings is to train word embeddings for each language independently using monolingual resources and then aligning them using a linear transformation (Artetxe et al., 2018). Approaches for word embedding alignment can be grouped into three categories: supervised (Mikolov et al., 2013; Lazaridou et al., 2015), semisupervised (Artetxe et al., 2017) and unsupervised (Artetxe et al., 2018; Alvarez-Melis and Jaakkola, 2018). Supervised approaches require a bilingual dictionary with a few thousand entries to learn the mapping. Semisupervised procedures need a small seed dictionary. Unsupervised approaches can align word embeddings without any bilingual data but, as shown by Vulić et al. (2019), they are only effective when the two languages are similar enough, restricting their applicability.

In this use case, we use our resource as the initial seed dictionary for semisupervised alignment of word embeddings for language pairs where unsupervised methods fail. We select three such language pairs – English/Japanese, English/Chinese and English/Tamil – and show that VecMap,¹³ a semisupervised method, can successfully employ our NE resource to align these languages. VecMap implements the method proposed by Artetxe et al. (2018), which is a state-of-the-art method for unsupervised cross-lingual word embedding mapping. It creates an initial set of word pairings based

¹³<https://github.com/artetxem/vecmap>

	Eng-Jpn	Eng-Tam	Eng-Zho
Unsupervised	0.0	0.0	0.0
Semisupervised	30.43	14.4	30.1

Table 6: P@1 BLI results with unsupervised VecMap compared to semisupervised VecMap, which uses our NE resource for initialization

on the distribution of words in their similarity matrix. Then it employs a self-learning method to improve the mapping iteratively.

We evaluate the embeddings on the Bilingual Lexicon Induction (BLI) task and the gold dataset provided by MUSE (Conneau et al., 2018). We use Wikipedia fast-Text embeddings (Bojanowski et al., 2017) as monolingual input vectors and report precision at one (P@1) for the unsupervised and semisupervised approaches in Table 6. While the fully unsupervised method fails to align these languages, the semisupervised approach based on our resource has much better results confirming that our NE resource can be effectively used as seed data.

6. Resource

We release a resource of named entities for 1340 languages, 1134 of which are lowest-resource.¹⁴ The resource mainly contains people and location NEs. The total number of NEs is 674,493, so there are 503 NEs per language on average with at least 300 NEs in 95% of the languages. The three best represented language families (Dryer and Haspelmath, 2013) are Austronesian, Niger-Congo and Indo-European. However, our coverage broadly includes all major areas of linguistic diversity, including Amazonian (e.g., Kaingang), African (e.g., Sango) and Papua New Guinea (e.g., Saniyo-Hiyewe).

7. Conclusion

We presented CLC-BN, a new method that identifies named entity correspondences and trains a neural transliteration model on them. CLC-BN does not need any other bilingual resources beyond the parallel corpus nor a word aligner or seed data. We showed that it outperforms prior work on silver data and human-annotated gold data. We created a new NE resource for 1340 languages by applying CLC-BN to the Parallel Bible Corpus and illustrated its utility by demonstrating good performance on two downstream tasks: knowledge graph augmentation and bilingual lexicon induction.

Acknowledgments. This work was funded by the European Research Council (grant #740516) and the German Federal Ministry of Education and Research (BMBF, grant #01IS18036A).

8. Bibliographical References

Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., and Vollgraf, R. (2019). Flair: An easy-to-use framework for state-of-the-art nlp. In *NAACL*

¹⁴Our NEs resource is freely available at http://cistern.cis.lmu.de/ne_bible/

- 2019, *2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59.
- Alvarez-Melis, D. and Jaakkola, T. S. (2018). Gromov-wasserstein alignment of word embedding spaces. *arXiv preprint arXiv:1809.00013*.
- Artetxe, M., Labaka, G., and Agirre, E. (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.
- Artetxe, M., Labaka, G., and Agirre, E. (2018). A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798, Melbourne, Australia, July. Association for Computational Linguistics.
- Bird, S. (2020). Decolonising speech and language technology. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519.
- Brown, P. F., Della Pietra, S. A., Della Pietra, V. J., and Mercer, R. L. (1993). The mathematics of statistical machine translation: Parameter estimation. *Computational linguistics*, 19(2):263–311.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, October. Association for Computational Linguistics.
- Cohen, W., Ravikumar, P., and Fienberg, S. (2003). A comparison of string metrics for matching names and records. In *Kdd workshop on data cleaning and object consolidation*, volume 3, pages 73–78.
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Conneau, A. and Lample, G. (2019). Cross-lingual language model pretraining. *Advances in neural information processing systems*, 32.
- Dasigi, P. and Diab, M. (2011). Named entity transliteration generation leveraging statistical machine translation technology. In *Proceedings of the 3rd Named Entities Workshop (NEWS 2011)*, pages 106–111.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019a). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019b). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Imani, A., Sabet, M., Keremşenel, L., Dufter, P., Yvon, F., and Schütze, H. (2021). Graph algorithms for multiparallel word alignment. In *The 2021 Conference on Empirical Methods in Natural Language Processing*.
- Jalili Sabet, M., Dufter, P., Yvon, F., and Schütze, H. (2020). SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online, November. Association for Computational Linguistics.
- Jaro, M. A. (1989). Advances in record-linkage methodology as applied to matching the 1985 census of tampa, florida. *Journal of the American Statistical Association*, 84(406):414–420.
- Joshi, P., Santy, S., Budhiraja, A., Bali, K., and Choudhury, M. (2020). The state and fate of linguistic diversity and inclusion in the nlp world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293.
- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., et al. (2007). Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th annual meeting of the association for computational linguistics companion volume proceedings of the demo and poster sessions*, pages 177–180.
- Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., and Dyer, C. (2016). Neural architectures for named entity recognition. In *Proceedings of NAACL-HLT*, pages 260–270.
- Landis, J. R. and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Lazaridou, A., Dinu, G., and Baroni, M. (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.
- Li, J., Sun, A., Han, J., and Li, C. (2020). A survey on deep learning for named entity recognition. *IEEE Transactions on Knowledge and Data Engineering*.
- Li, B., He, Y., and Xu, W. (2021). Cross-lingual named entity recognition using parallel corpus: A new approach using xlm-roberta alignment. *arXiv preprint arXiv:2101.11112*.
- Liang, P., Taskar, B., and Klein, D. (2006). Alignment by agreement. In *Proceedings of the Human Lan-*

- guage Technology Conference of the NAACL, Main Conference, pages 104–111.
- Luong, M.-T., Pham, H., and Manning, C. D. (2015). Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Ma, X. and Hovy, E. (2016). End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074.
- Mikolov, T., Le, Q. V., and Sutskever, I. (2013). Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Ngo-Ho, A.-K. and Yvon, F. (2019). Neural baselines for word alignments. In *International Workshop on Spoken Language Translation*.
- Ni, J., Dinu, G., and Florian, R. (2017). Weakly supervised cross-lingual named entity recognition via effective annotation and representation projection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1470–1480.
- Och, F. J. and Ney, H. (2000). Improved statistical alignment models. In *Proceedings of the 38th annual meeting of the association for computational linguistics*, pages 440–447.
- Och, F. J. and Ney, H. (2003). A systematic comparison of various statistical alignment models. *Computational linguistics*, 29(1):19–51.
- Östling, R., Tiedemann, J., et al. (2016). Efficient word alignment with markov chain monte carlo. *The Prague Bulletin of Mathematical Linguistics*.
- Prabhakar, D. K. and Pal, S. (2018). Machine transliteration and transliterated text retrieval: a survey. *Sādhanā*, 43(6):1–25.
- Rios, A., Amrhein, C., Aepli, N., and Sennrich, R. (2021). On biasing transformer attention towards monotonicity. In *2021 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Sabet, M. J., Dufter, P., Yvon, F., and Schütze, H. (2020). Simalign: High quality word alignments without parallel training data using static and contextualized embeddings. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643.
- Sajjad, H., Fraser, A., and Schmid, H. (2011). An algorithm for unsupervised transliteration mining with an application to word alignment. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 430–439.
- Semmar, N. and Saadane, H. (2013). Using transliteration of proper names from Arabic to Latin script to improve English-Arabic word alignment. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1022–1026, Nagoya, Japan, October. Asian Federation of Natural Language Processing.
- Severini, S., Hangya, V., Fraser, A., and Schütze, H. (2020). Combining word embeddings with bilingual orthography embeddings for bilingual dictionary induction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6044–6055, Barcelona, Spain (Online), December. International Committee on Computational Linguistics.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Tsai, C.-T., Mayhew, S., and Roth, D. (2016). Cross-lingual named entity recognition via wikification. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 219–228.
- Vulić, I., Glavaš, G., Reichart, R., and Korhonen, A. (2019). Do we really need fully unsupervised cross-lingual embeddings? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4407–4418, Hong Kong, China, November. Association for Computational Linguistics.
- Wang, M. and Manning, C. D. (2014). Cross-lingual projected expectation regularization for weakly supervised learning. *Transactions of the Association for Computational Linguistics*, 2:55–66.
- Wang, Z., Lv, Q., Lan, X., and Zhang, Y. (2018). Cross-lingual knowledge graph alignment via graph convolutional networks. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 349–357.
- Wu, S. and Cotterell, R. (2019). Exact hard monotonic attention for character-level transduction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1530–1537.
- Wu, W. and Yarowsky, D. (2018). A comparative study of extremely low-resource transliteration of the world’s languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Wu, Q., Lin, Z., Wang, G., Chen, H., Karlsson, B. F., Huang, B., and Lin, C.-Y. (2020). Enhanced meta-learning for cross-lingual named entity recognition with minimal resources. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9274–9281.
- Wu, Q., Lin, Z., Karlsson, B. F., Huang, B., and Lou, J.-G. (2021). Unitrans: unifying model transfer and data transfer for cross-lingual named entity recognition with unlabeled data. In *Proceedings of the*

Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence, pages 3926–3932.

Yadav, V. and Bethard, S. (2018). A survey on recent advances in named entity recognition from deep learning models. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2145–2158.

9. Language Resource References

Agić, Ž. and Vulić, I. (2019). Jw300: A wide-coverage parallel corpus for low-resource languages. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3204–3210.

Al-Rfou, R., Kulkarni, V., Perozzi, B., and Skiena, S. (2015). Polyglot-ner: Massive multilingual named entity recognition. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 586–594. SIAM.

Benites, Fernando and Duivesteijn, Gilbert François and von Däniken, Pius and Cieliebak, Mark. (2020). *TRANSLIT: A Large-scale Name Transliteration Resource*. European Language Resources Association.

Bojanowski, Piotr and Grave, Edouard and Joulin, Armand and Mikolov, Tomas. (2017). *Enriching Word Vectors with Subword Information*.

Alexis Conneau and Guillaume Lample and Marc’Aurelio Ranzato and Ludovic Denoyer and Herv’e J’egou. (2018). *Word Translation Without Parallel Data*.

Matthew S. Dryer and Martin Haspelmath. (2013). *WALS Online*. Max Planck Institute for Evolutionary Anthropology.

Ehrmann, M., Turchi, M., and Steinberger, R. (2011). Building a multilingual named entity-annotated corpus using annotation projection. In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 118–124.

Hermjakob, Ulf and May, Jonathan and Knight, Kevin. (2018). *Out-of-the-box universal romanization tool uroman*.

Mayer, Thomas and Cysouw, Michael. (2014). *Creating a massively parallel bible corpus*.

Navigli, Roberto and Ponzetto, Simone Paolo. (2012). *BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network*. Elsevier.

Pan, Xiaoman and Zhang, Boliang and May, Jonathan and Nothman, Joel and Knight, Kevin and Ji, Heng. (2017). *Cross-lingual name tagging and linking for 282 languages*.

Wu, Winston and Vyas, Nidhi and Yarowsky, David. (2018). *Creating a translation matrix of the Bible’s names across 591 languages*.

A. Reproducibility Information

We run our method on up to 48 cores of Intel(R) Xeon(R) CPU E7-8857 v2 with 1TB memory and a single GeForce GTX 1080 GPU with 8GB memory. CLC-BN is implemented in Python and takes approximately 2 minutes to run for one language. The neural model is implemented in PyTorch and has one encoder and one decoder layer (batch size 16, hidden layer size 32, learning rate 0.01, dropout 0.4, 24K parameters). We use Luong et al. (2015)’s attention. Each training of the neural transliteration model requires at most 10 minutes. SimAlign (Sabet et al., 2020) alignments are obtained using multilingual BERT (Devlin et al., 2019b). We use subword alignments and the forward alignment to ensure that all English NEs are aligned. Eflomal (Östling et al., 2016) alignments are obtained with default parameters and the forward alignment. The Jaro distance is calculated using the Python library `textdistance`.¹⁵ For the cross-lingual word alignment experiment we used the latest VecMap code available in its git repository¹⁶ (no snapshot is available). We ran it using the `< --unsupervised >` and `< --semi_supervised >` switches. All other parameters are left as their default value. The monolingual word alignments are downloaded from fastText’s official website.¹⁷

¹⁵<https://pypi.org/project/textdistance/>

¹⁶commit ID: b82246f6c249633039f67fa6156e51d852bd73a3

¹⁷<https://fasttext.cc/docs/en/pretrained-vectors.html>