Weakly Supervised POS Taggers Perform Poorly on *Truly* Low-Resource Languages

Katharina Kann*
Center for Data Science
New York University, USA
kann@nyu.edu

Ophélie Lacroix*
Siteimprove, Denmark
ola@siteimprove.com

Anders Søgaard
Department of Computer Science
University of Copenhagen, Denmark
soegaard@di.ku.dk

Abstract

Part-of-speech (POS) taggers for low-resource languages which are exclusively based on various forms of weak supervision – e.g., cross-lingual transfer, type-level supervision, or a combination thereof – have been reported to perform almost as well as supervised ones. However, weakly supervised POS taggers are commonly only evaluated on languages that are very different from *truly* low-resource languages, and the taggers use sources of information, like high-coverage and almost error-free dictionaries, which are likely not available for resource-poor languages. We train and evaluate state-of-the-art weakly supervised POS taggers for a typologically diverse set of 15 truly low-resource languages. On these languages, given a realistic amount of resources, even our best model gets only less than half of the words right. Our results highlight the need for new and different approaches to POS tagging for truly low-resource languages.

Introduction

Part-of-speech (POS) tagging can be very helpful in many natural language processing (NLP) applications, including machine translation, question answering, or relation extraction, especially for low-resource languages (Sennrich and Haddow 2016; Nadejde et al. 2017; Currey and Heafield 2019). POS taggers assign a syntactic category, i.e., a part of speech, to each token in a sentence, thus providing a rudimentary syntactic analysis of the sentence. While ambiguity and unseen words make POS tagging non-trivial, supervised POS tagging is often considered a relatively simple problem, reaching a performance close to inter-annotator agreement (Bohnet et al. 2018, 97.96% accuracy for English). The fact that weakly supervised POS taggers sometimes obtain almost-as-good scores makes POS tagging come across as a task where there is little room for improvement. We show that on the contrary, there is *a lot* of room for improvement.

The accuracy of POS taggers naturally depends on the quality of available training data, as well as the granularity of the POS inventory, i.e., the size of the tag set. Nevertheless, for almost all resource-rich languages considered in

the literature, supervised POS taggers have reported tagging accuracies in the high 90s.¹ However, supervised POS taggers typically require corpora that have been manually created by professional linguistic annotators.² In the absence of such corpora, we have to resort to alternative sources of (weak) supervision. One very popular source of such supervision is tag dictionaries (Li, Graça, and Taskar 2012; Garrette and Baldridge 2013; Täckström et al. 2013; Wisniewski et al. 2014; Moore 2015; Plank and Agić 2018).

In this paper, we revisit and critically examine existing methods for weakly supervised POS tagging. While such approaches have been created to meet the needs of low-resource languages, many experiments with weakly supervised POS taggers have been limited to languages for which good resources actually exist. This has practical reasons: it is easier to obtain translations, dictionaries, and benchmark data for widely studied languages. However, since the resource-rich languages are *not* a representative sample of the world’s languages, this means we have no guarantee that results scale to truly low-resource languages.

Contributions We study the performance of weakly supervised POS taggers for actual low-resource languages. To this end, we consider a language *truly low-resource* for a task when no or almost no resources, i.e., annotated corpora or manually created dictionaries, are available for that task in that language. Assuming a setting in which we have access to no annotated data in the target language at all, we investigate several state-of-the-art baselines, as well as two cross-lingual transfer variants of a state-of-the-art model for the high-resource case (Plank, Søgaard, and Goldberg 2016), extended to a multi-task learning (MTL) architecture in order to provide additional character-level supervision (Kann et al. 2018). Our results show that POS tagging is *still a difficult problem in the truly low-resource setting*: **Our strongest baseline** obtains a macro-average accuracy of **39.11%** over 15 languages, even though it was devel-

¹The CoNLL 2018 shared task (Zeman et al. 2018) on universal dependencies reported macro-average F_1 scores of up to 96.23% on a set of 61 treebanks in high-resource languages.

²See Hovy et al. (2015) for experiments with crowdsourcing POS annotation, indicating that sometimes, lay annotators can be used at a small cost in accuracy.

*The first two authors contributed equally.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

oped specifically for low-resource languages. **Our best architecture**, based on cross-lingual transfer and character-level MTL, is slightly better, with a macro-average accuracy of 42.25%, but our experiments emphasize the need for additional research on POS tagging for truly low-resource languages. We provide our preprocessed data for all languages under *anonymized* in order to facilitate future research.

Disclaimer In this study, we do not consider the case where small amounts of labeled data is (also) available for the target language. While it is easy to think that annotating a few hundred sentences should be cheap, developing guidelines, as well as finding and training annotators for truly low-resource languages is challenging and often impossible. Moreover, we do not look at the potential gains from using contextualized word embeddings from pretrained language models (Peters et al. 2018; Devlin et al. 2019), since even high-quality language models can be hard to train for truly low-resource languages, because data is scarce and exhibits significant spelling variation, and because language detection software has limited precision for truly low-resource languages (Baldwin and Lui 2010).

Unsupervised and Weakly Supervised POS Tagging

Since the resources required to train supervised taggers are expensive to create and unlikely to exist for the majority of the world’s languages, many unsupervised and weakly supervised methods—the latter often building on the former—have been developed. We now briefly survey this work.

Unsupervised POS tagging Hidden Markov models (HMMs) and their variants have been standard models for POS induction (Merialdo 1994; Goldwater and Griffiths 2007). Christodoulopoulos, Goldwater, and Steedman (2010) evaluated seven different POS induction systems, from the well-known clustering method of Brown et al. (1992) to Berg-Kirkpatrick et al. (2010), who incorporated linguistically motivated features into HMMs. Subsequently, Christodoulopoulos, Goldwater, and Steedman (2011) proposed a Bayesian multinomial mixture model instead of a sequence model to infer clusters from unlabeled text. More recent studies include (Stratos, Collins, and Hsu 2016) who proposes to use anchor HMMs.

POS tags could then be assigned to clusters, e.g., by using monolingual dictionaries, however, in most studies gold annotated data are used to find the best match of POS-tags, hence reporting upper-bound performance of their systems. While all approaches mentioned above were reported to perform reasonably, none of the results were obtained for actual low-resource languages (and, in a low-resource setting such as when using low-resource dictionaries), e.g., 75.5% for English (Berg-Kirkpatrick et al. 2010), 74.4% for German, and 61.5% for Arabic (Christodoulopoulos, Goldwater, and Steedman 2011).

Weakly supervised POS tagging Weakly supervised methods for POS tagging tend to rely on the projection of information across word alignments between parallel corpora (Täckström et al. 2013; Agić, Hovy, and Søgaard 2015; Plank and Agić 2018), cross-lingual features obtained from

parallel data, comparable data, or seed bilingual dictionaries (Gouws and Søgaard 2015; Fang and Cohn 2017), or monolingual tag dictionaries that constrain the search space of unsupervised POS induction (Li, Graça, and Taskar 2012; Wisniewski et al. 2014; Garrette and Baldrige 2013). Strategies using dictionaries extracted from WIKTIONARY have been popular for learning from noisily labeled sentences (Li, Graça, and Taskar 2012; Wisniewski et al. 2014). Li, Graça, and Taskar (2012) proposed a feature-based maximum entropy emission model for learning POS taggers from texts which have been labeled using monolingual dictionaries. However, they evaluate the performance of their system exclusively on resource-rich languages for which such dictionaries cover a large amount of their vocabulary, resulting in performances over 80%. Garrette and Baldrige (2013) proposed a maximum entropy Markov model for learning POS taggers from noisily labeled sentences. They labeled a raw corpus with the help of a manually annotated dictionary that was expanded by label propagation. Notably, the authors did evaluate on actual low-resource languages. However, to actually achieve reasonable results, they bootstrapped their model with a small set of manually annotated sentences.

Our work is also related to Fang and Cohn (2017), in that we do not rely on parallel text, since we are interested in a setting in which large amounts of parallel corpora are not available, as might be the case for true low-resource languages. However, in contrast, we do not assume a small amount of training data for the target language. Further, our approach is similar to Täckström et al. (2013) in creating silver-standard training data from type-level dictionaries, but we make use of the state-of-the-art architecture by Plank, Søgaard, and Goldberg (2016) instead of employing HMMs. Our study is also similar to Agić, Hovy, and Søgaard (2015) in evaluating approaches to weakly supervised POS tagging across many low-resource languages, but differs in that we are interested in methods without need for parallel corpora.

Model Architecture

Hierarchical POS tagging long short-term memory networks (LSTMs), such as the architecture proposed by Plank, Søgaard, and Goldberg (2016), receive both word-level and subword-level input. They perform well, even on unseen words, due to their ability to associate subword-level patterns with POS tags. This is important in a low-resource setting. However, hierarchical LSTMs are also very expressive, and thus prone to overfitting. In order to overcome this, we additionally train our models on subword-level auxiliary tasks (Kann et al. 2018) to regularize the character-level encoding in hierarchical LSTMs. Such a model is still able to make predictions about unknown words, but the subword-level auxiliary task should prevent it from overfitting.

Hierarchical LSTMs with Character-Level Decoding

For the hierarchical sequence labeling LSTM, we follow Plank, Søgaard, and Goldberg (2016): our subword-level LSTM is bidirectional and operates on the character level (Ballesteros, Dyer, and Smith 2015). Its input is the char-

acter sequence of each input word, represented by the embedding sequence $c_1, c_2, \dots, c_m$. The final character-based representation of each word is the concatenation of the two last LSTM hidden states:

$$v_{c,i} = \text{conc}(\text{LSTM}_{c,f}(c_{1:m}), \text{LSTM}_{c,b}(c_{m:1})) \quad (1)$$

Here, $\text{LSTM}_{c,f}$ and $\text{LSTM}_{c,b}$ denote a forward and backward LSTM, respectively.

Second, a context bi-LSTM operates on the word level. Like Plank, Sogaard, and Goldberg (2016), we use the term “context bi-LSTM” to denote a bidirectional LSTM which, in order to generate a representation for input element i , encodes all elements up to position i with a forward LSTM and all elements from n to i using a backward LSTM. For each sentence represented by embeddings $w_1, w_2, \dots, w_n$, its input are the concatenation of the word embeddings with the outputs of the subword-level LSTM: $\text{conc}(w_1, v_{c,1}), \text{conc}(w_2, v_{c,2}) \dots, \text{conc}(w_n, v_{c,n})$. The final representation which gets forwarded to the next part of the network is again the concatenation of the last two hidden LSTM states:

$$v_{w,i} = \text{conc}(\text{LSTM}_{w,f}(\text{conc}(w, v_c)_{1:i}), \text{LSTM}_{w,b}(\text{conc}(w, v_c)_{n:i})) \quad (2)$$

This is then passed on to a classification layer.

Character-Based Seq2Seq Model

The second component of our model is based on a character-level sequence-to-sequence (seq2seq) architecture. It consists of a bi-LSTM encoder which is connected to an LSTM decoder (Sutskever, Vinyals, and Le 2014).

Encoding The encoder is the character-level bi-LSTM described above, and, thus, yields the representation

$$v_{c,i} = \text{conc}(\text{LSTM}_{c,f}(c_{1:m}), \text{LSTM}_{c,b}(c_{m:1})) \quad (3)$$

for an input word embedded as $c_1, c_2, \dots, c_m$. Parameters of the character-level LSTM are shared between the sequence labeling and the seq2seq part of our model.

Decoding The decoder receives the concatenation of the last hidden states, $v_{c,i}$, as input. In particular, we do not use an attention mechanism (Bahdanau, Cho, and Bengio 2015), since our goal is not to improve performance on the auxiliary task, but instead to encourage the encoder to learn better word representations. The decoder is trained to predict each output character y_t dependent on $v_{c,i}$ and previous predictions $y_1, \dots, y_{t-1}$ as

$$p(y_t | \{y_1, \dots, y_{t-1}\}, v_{c,i}) = g(y_{t-1}, s_t, v_{c,i}) \quad (4)$$

for a non-linear function g and the LSTM hidden state s_t . The final softmax output layer is calculated over the character vocabulary of the language.

Multi-Task Learning

The character-level LSTM is shared between the sequence labeling and the seq2seq components of our network. All model parameters, including all embeddings, are updated during training.

We want to train our neural model jointly on (i) a low-resource main task, i.e., POS tagging, and (ii) a character-level auxiliary task, namely word autoencoding. Therefore, we want to maximize the following joint log-likelihood:

$$\mathcal{L}(\theta) = \sum_{(l,s) \in \mathcal{D}_{\text{POS}}} \log p_{\theta}(l | s) + \sum_{(in,out) \in \mathcal{D}_{\text{aux}}} \log p_{\theta}(out | in) \quad (5)$$

Here, $\mathcal{D}_{\text{POS}}$ denotes the POS tagging training data, with s being the input sentence and l the corresponding label sequence. $\mathcal{D}_{\text{aux}}$ is our auxiliary task training data with examples consisting of input in and output out . The set of model parameters θ is the union of the set of parameters of the sequence labeling and the seq2seq part. Parameters of the character-level LSTM are shared between tasks.

Word autoencoding (AE) Our auxiliary task consists of reproducing a given input character sequence in the output. Thus, training examples are of the form $w \mapsto w$, where $w \in \mathcal{V}_L$ for the vocabulary $\mathcal{V}_L$ of L . Word autoencoding has been used as an auxiliary task before, e.g., by Vosoughi, Vijayaraghavan, and Roy (2016).

Weak and Cross-Lingual Supervision

The model described above relies on full supervision. However, in our setting, we need to rely on alternative resources, including raw corpora and linguistic resources, and cross-lingual transfer to create training examples. The raw corpora which we leverage for our approaches consist of cleaned WIKIPEDIA texts for our languages. Preprocessing involves (a) segmentation of text into sentences; (b) tokenization; (c) removal of sentences that include at least one of (i) mostly foreign characters; (ii) mostly symbols or punctuation; (iii) no words from our dictionaries. Most of the texts that we extract for our low-resource languages are predominantly made up of ambiguous sentences (i.e., most words in these sentences are mapped to more than one POS tag in our dictionaries) which makes it impossible to extract, for training, only sentences that would be fully (and unambiguously) tagged with our dictionaries. Numbers of sentences and tokens extracted for each language are shown in Table 1.

In order to obtain silver-standard training data for POS tagging in our low-resource languages, we rely on cross-lingual transfer. In particular, we assume the following to be available for each low-resource language: (i) raw text, e.g., the WIKIPEDIA corpora we just described; (ii) a bilingual dictionary B , containing translations which consist of pairs of words (w_l, w_h) in the low-resource language L_l and a high-resource language L_h , respectively; (iii) large amounts of gold POS-annotated data D in the high-resource language L_h ; and, optionally, (iv) a monolingual tag dictionary M . Given those resources, we propose two cross-lingual transfer methods, which we will outlay in the following.

Frequency-based annotation (FREQ) The high-level idea of our first approach is to tag each token with the POS tag which has been most frequently assigned to its high-

Language		Treebank Data (test)		UNIMORPH	WIKIDATA+PANLEX	WIKIPEDIA (# tagged)		Embeddings
code	family	sentences	tokens	entries	translations	sentences	tokens	entries
am	AA	1,095	10k	-	2.7k	777	17.9k	10k
be	IE	68	1.3k	-	35.3k	7,385	101.9k	93k
br	IE	888	10.3k	-	12.2k	9,083	112.9k	39k
fo	IE	1,208	10.0k	45.4k	2.9k	9,958	144.6k	12k
hsb	IE	623	10.7k	-	4.6k	1,858	30.2k	10k
hy	IE	514	11.4k	338k	65.1k	3,560	71.4k	47k
kmr	IE	734	10.1k	-	4.6k	3,225	48.3k	24k
lt	IE	55	1.0k	34.1k	38.9k	11,464	117.2k	100k
mr	IE	47	0.4k	-	23.4k	4,886	55.2k	47k
mt	AA	100	2.3k	-	2.1k	2,361	43.9k	16k
bxr	Mo	908	10.0k	-	2.7k	2,308	37.8k	28k
kk	Tu	1,047	10.1k	-	63.5k	12,273	122.4k	100k
ta	Dr	120	2.2k	-	27.1k	5,772	76.2k	100k
te	Dr	146	0.7k	-	28.0k	7,872	90.9k	100k
tl	Au	55	0.2k	-	6.8k	5,871	97.6k	41k
<hr/>								
de	IE	1,000	21.3k	179.3k	90.2k	12,162	195.1k	100k
es	IE	1,000	23.3k	382.9k	59.7k	15,209	276.6k	100k
it	IE	1,000	23.7k	509.5k	59.7k	10,254	170.0k	100k
pt	IE	1,000	23.4k	303.9k	47.9k	12,674	195.2k	100k
sv	IE	1,000	19.1k	78.4k	58.8k	10,243	134.5k	100k

Table 1: Resources for our low-resource languages (up) and high-resource languages (down). Language families: Afro-Asiatic (AA), Austronesian (Au), Dravidian (Dr), Indo-European (IE), Mongolic (Mo), and Turkic (Tu).

resource language translations:

$$\begin{aligned}
\text{TAG}(w_l) &= t \in \cup_{w_h \in \text{Tr}(w_l)} \text{POS}(w_h) : \\
&\quad \text{freq}_{\text{Tr}(w_l)}(t) \geq \text{freq}_{\text{Tr}(w_l)}(s) \\
&\quad \forall s \in \cup_{w_h \in \text{Tr}(w_l)} \text{POS}(w_h)
\end{aligned} \tag{6}$$

Here, $\text{Tr}(w_l)$ denotes all possible translations of w_l in B , $\text{POS}(w_h)$ is the set of all attested POS tags of w_h in D and $\text{freq}_{\text{Tr}(w_l)}(x)$ is the number of times tag x has been assigned to any word in $\text{Tr}(w_l)$ in D . B is a type-level resource, while the training data D in the high-resource language is token-based, i.e., words are tagged in context.

A word which appears neither in B nor in D is not tagged. We switch learning off for those tokens, i.e., we mask the calculation of the loss and, thus, do not take them into account during optimization. Raw sentences which do not contain at least one tagged word are discarded.

Ambiguous annotation (AMB) We next propose to annotate our raw sentences with *noisy* or *ambiguous* labels. We tag each token with all POS tags we consider possible, given dictionary B and high-resource language data D , i.e., we assign to a token w_l all tags in $\text{POS}(w_l)$ with

$$\text{POS}(w_l) = \cup_{w_h \in \text{Tr}(w_l)} \text{POS}(w_h), \tag{7}$$

and all variables denoting the same quantities as before. In order to include our monolingual dictionaries, which give a set of possible tags $\text{POS}_M(w_l)$ for each word w_l in the low-resource language, we extend this to

$$\text{POS}(w_l) = \cup_{w_h \in \text{Tr}(w_l)} \text{POS}(w_h) \cup \text{POS}_M(w_l) \tag{8}$$

Words that do not appear in either M , B or D are tagged with all possible tags. As before, we discard raw sentences which contain only tokens without information.

For training, we adapt the calculation of the cross entropy to ambiguous annotation. In particular, we obtain the loss of each example by treating the tag in $\text{POS}(w_l)$ which obtains the highest probability under our model as the gold label.

Languages and Resources

The goal of this work is to evaluate state-of-the-art unsupervised or weakly supervised POS-tagging strategies on truly low-resource languages. For this, we select 15 languages from the Universal Dependencies (UD) project v2.1 (Nivre et al. 2017).³ All chosen languages are low-resource languages: 10 of them (Belarusian (be), Buryat (bxr), Upper Sorbian (hsb), Armenian (hy), Kazakh (kk), Kurmanji (kmr), Lithuanian (lt), Marathi (mr), Tamil (ta) and Telugu (te)) have less than 10k tokens for training in UD, and the other 5 languages (Amharic (am), Breton (br), Faroese (fo), Maltese (mt) and Tagalog (tl)) have no training data at all. Note, however, that we consider the training sets only to determine languages which are low-resource languages, and we do not make use of *any* training data in our experiments. Our languages represent 6 different languages families: Afro-Asiatic (am and mt), Austronesian (tl), Dravidian (ta and te), Indo-European (be, br, fo, hsb, hy, kmr, lt, and mr), Mongolic (bxr), and Turkic (kk); cf. Table 1.

English is used as the high-resource language for cross-lingual transfer in our experiments. Note that a more informed choice of source language might be able to obtain better results, but that bilingual dictionaries are more likely to exist from/to English.

³<http://universaldependencies.org/>

Dictionaries

Bilingual dictionaries A possible resource for weakly supervised approaches to cross-lingual POS tagging are bilingual dictionaries that contain word-to-word translations. They can be used for transferring information from resource-rich languages to low-resource languages, either by replacing words directly, transferring statistics, or for inducing cross-lingual word representations as done, e.g., by Faruqui and Dyer (2014). Some bilingual dictionaries can be downloaded from the WIKTIONARY user page of Matthias Buchmeier⁴ or from the WIKT2DICT project website⁵. Unfortunately, those dictionaries do not include many low-resource languages (only 1 (lt) in our 15 languages is covered). We therefore decide to rely on freely available resources for extracting bilingual dictionaries: PANLEX SWADESH⁶ and WIKIDATA⁷. The PANLEX SWADESH Corpora (Baldwin, Pool, and Colowick 2010) gathers lists of 317 words (and synonyms) for over 600 languages. From this, we retrieve small bilingual dictionaries for almost all low-resource languages we consider.⁸ Our second resource, WIKIDATA, is an online database collecting more than 28 million links to WIKIPEDIA and WIKTIONARY pages across different languages. The English pages align the multilingual sites, which enables us to extract bilingual dictionary entries. We extract bilingual dictionaries (English to target language) from these pages for our 15 low-resource languages.⁹ Statistics are provided in Table 1.

Monolingual dictionaries A common approach for obtaining monolingual tag dictionaries, i.e., lists of tokens matched with their possible POS tags, for low-resource languages is to extract this information from WIKTIONARY,¹⁰ a multilingual free online dictionary. However, as an online collaborative tool that can be edited by any user, it is noisy and prone to errors. This is particularly true for low-resource languages, which is one of the reasons why studies such as Li, Graça, and Taskar (2012) and Wisniewski et al. (2014) limited themselves to a few (not truly resource-poor) languages.

The UNIMORPH project (Kirov et al. 2016), in contrast, was professionally created and contains morpho-syntactic information for 107 languages. However, this resource is only available for 3 of our languages (see Table 1), and POS tags are restricted to nouns, verbs, adjectives, as well as a small number of adverbs. We convert the UNIMORPH tagging scheme into UD POS tags for all languages.

Embeddings

We leverage monolingual word embeddings to improve our ability to generalize to unseen words in the target language. This can be seen as MTL with a language modeling auxiliary task. However, since the use of word embeddings is de facto

standard in NLP, we consider them a basic part of our model. We employ the POLYGLOT embeddings which have been built from WIKIPEDIA texts and made available by Al-Rfou, Perozzi, and Skiena (2013).¹¹ See Table 1 (last column) for the number of words covered by those embeddings.

Experiments

We explore a set of state-of-the-art unsupervised and weakly supervised POS tagging strategies. Note that we do not experiment with projected information from parallel corpora like Agić, Hovy, and Søgaaard (2015), due to datasets only being available for 7 out of our 15 languages. As mentioned in the disclaimer above, we also do not consider limited supervision or contextualized word embeddings.

CHR11 We employ the fully unsupervised method of Christodoulopoulos, Goldwater, and Steedman (2011)¹² (cf. related work section) as our first baseline. We learn clusters on raw text extracted from WIKIPEDIA using their Bayesian multinomial mixture model. To infer POS tags, we use our monolingual dictionaries to match each cluster with the POS tag that is most frequently associated with its tokens.

GAR13 Second, we compare to Garrette and Baldrige (2013)’s system¹³ (cf. related work section), which consists of a maximum entropy Markov model which learns from a corpus of noisily labeled sentences. For this, we make use of the raw WIKIPEDIA texts and our monolingual dictionaries.

PLA16 Plank, Søgaaard, and Goldberg (2016) is a hierarchical LSTM POS tagger relying on a combination of character embeddings and word embeddings. Plank, Søgaaard, and Goldberg (2016) train this architecture in a multi-task fashion, using the task of predicting the log frequency of the next word as an auxiliary task. We compute log frequencies for each word w_n as $\text{int}(\log(\text{freq}_{\text{train}}(w_{n+1})))$. We then learn POS taggers in a MTL setup: the main task learns to predict POS tags using the ambiguous learning strategy (data is created the same way as for AMB), and the auxiliary task learns a simplified language model. The auxiliary task training data consists of sentences from WIKIPEDIA.

We mostly adopt the hyperparameters from Plank, Søgaaard, and Goldberg (2016). The number of dimensions of our word embeddings is 64 (which is also the dimension of the POLYGLOT embeddings). We use one hidden layer of dimension 100 for both the word and the character LSTM. In our use of dropout we also follow Plank, Søgaaard, and Goldberg (2016). However, we further add character dropout with a coefficient of 0.25 to improve regularization. We train a minimum and maximum number of 15 and 30 epochs, respectively, and terminate if no improvement in the training loss is detected for 3 consecutive epochs. At test time, we use the model which obtained the lowest loss.

We compute global accuracy scores over POS tags for all systems. Our evaluation is token-based, and each reported result is an average over 5 runs with different random seeds.

⁴https://en.wiktionary.org/wiki/User:Matthias_Buchmeier

⁵<https://github.com/juditacs/wikt2dict>

⁶<https://panlex.org>

⁷<https://www.wikidata.org>

⁸PANLEX SWADESH data is not available for kmr and mr.

⁹We use the Mongolian and the Kurdish version for, respectively, Buryat and Kurmanji, which are closely related languages.

¹⁰<https://www.wiktionary.org>

¹¹<https://sites.google.com/site/rmyeid/projects/polyglot>

¹²<https://github.com/christos-c/bmmm>

¹³<https://github.com/dhgarrette/low-resource-pos-tagging-2014>

Results

The results of all experiments are presented in Table 2; the first three columns show the baseline results, the last four columns contain our proposed approaches. Our most important observation is that none of the approaches perform well on these languages. In fact, on average, no method gets even half of the tags right. This shows there is still a lot of room for exploring new and radically different approaches to learning POS taggers for (truly) low-resource languages. In addition, we make the following four observations: (i) We see that the unsupervised baseline, CHR11, obtains the lowest accuracy on average over all languages. This shows that, while all POS taggers are poor, distant supervision from tag dictionaries may provide some signal. (ii) On average over all languages, both single-task cross-lingual neural approaches, i.e., AMB and FREQ, outperform all baselines. Further, the difference between AMB and FREQ is with 0.0032 small. Thus, it seems that both cross-lingual strategies, i.e., ambiguous labeling and using the most frequent tag as the gold label, work similarly well on average. Looking at individual languages, however, up to around 0.1 performance difference can be found for am, hsb, hy, kmr, and te. Most likely, this difference might be explained by the frequency difference between the most frequent and other possible tags in the respective languages. (iii) Looking at PLA16, we find that a language modeling auxiliary task seems to hurt performance on average. PLA16 uses the ambiguous annotation strategy for cross-lingual transfer and should, thus, directly improve over AMB. However, it only performs better for three languages: am, ta, and tl. This contrast to Plank, Søggaard, and Goldberg (2016)’s results could be due to the small corpus sizes and, thus, mistakes in estimating word frequencies. (iv) The approaches with character-level MTL, i.e., AMB+AE and FREQ+AE, both improve over the respective single-task approaches. AMB+AE also obtains the highest average accuracy overall. Thus, in contrast to the language modeling auxiliary task, a character-level multi-task approach improves model performance in our setting.

Error Analysis

The most frequent POS tags in the test treebanks are NOUN (24.10% on average), PUNCT (15.46% on average), and VERB (13.57% on average). We suspect this to be the main reason why the cluster-based method CHR11 achieves competitive results: it tends to tag a huge part of the tokens with the most frequent POS tags, i.e., NOUN, PUNCT, and VERB. However, it fails on less frequent ones, losing against the other approaches. Bi-LSTMs tag more tokens belonging to less frequent categories correctly. However, a possible source of errors might be that the training sets we created for AMB and FREQ are fairly unbalanced: while the percentage of NOUN is 24.10% in the treebank test sets, only 15.19% of the tokens are assigned this tag for FREQ. The VERB tag, in contrast, which makes, on average, up for 13.57% of the tokens in the treebank test sets, is assigned to 24.65% of the words for FREQ. Similarly, the PUNCT tag is much more frequent for FREQ: 37.41%, as compared to 15.46% in the

test sets. For AMB, a total of 70.93% of the tokens have been tagged as possible NOUN. The next most frequent tags are VERB and PROP with 69.87% and, respectively, 61.34%. PUNCT is slightly less frequent than in the test sets: only 14.21% of the words have been tagged as such.¹⁴

Comparison with High-Resource Languages

Finally, the result that state-of-the-art taggers do not perform particularly well for the languages in our experiments is more valuable if we can find a plausible explanation for that. We, thus, repeat a subset of the main experiments for 5 high-resource languages: German (de), Spanish (es), Italian (it), Portuguese (pt) and Swedish (sv). We compare two different settings: First, we employ comparable resources to the previous experiment (-). Second, we use the dictionaries (named Wiki-ly) provided by Li, Graça, and Taskar (2012), which are extracted from WIKTIONARY and contain more reliable POS tags,¹⁵ and replace WIKIPEDIA texts with UD texts for training (+). We keep all hyperparameters the same as in the previous experiments. Again, we perform 5 training runs for all neural models and report average scores. We evaluate the scores on the PUD test sets (Hajič and Zeman 2017) of the UD treebanks for which there is a Wiki-ly dictionary available.

Results and discussion Results are shown in Table 3. Comparing the average results of the neural models with those for the low-resource languages, we find an improvement of 0.02 for AMB and 0.09 for FREQ, respectively. Thus, those approaches work slightly better for our high-resource languages, maybe because those are closer to English, our source language for cross-lingual transfer. However, when higher quality resources are used, tagging performance is even 0.31 higher for AMB. Similarly, CHR11 and GAR13 improve strongly over their results for the resource-poor languages. For FREQ, performance nearly doubles: it increases by 0.36. This highlights the importance of using high quality resources.

Related Work

POS tagging and other NLP sequence labeling tasks have been successfully approached using bidirectional LSTMs (Wang et al. 2015; Plank, Søggaard, and Goldberg 2016). Early work using such architectures relied on large annotated datasets, but Plank, Søggaard, and Goldberg (2016) showed that bi-LSTMs are not as reliant on data size as previously assumed. Their approach obtained state-of-the-art results for POS tagging in several languages, which is why we build upon it. Rei (2017) showed how an additional language modeling objective could improve performance for POS tagging. Neural networks make MTL via parameter sharing easy; thus, different task combinations have

¹⁴ Annotations for AMB are ambiguous, i.e., each token can have multiple tags, such that the percentages do not sum up to 1.

¹⁵ Note that Li, Graça, and Taskar (2012) use the Universal POS tags of Petrov, Das, and McDonald (2012), which do not exactly correspond to UD POS tags. We map, respectively, the CONJ and “.” tags to CCONJ and PUNCT. No PROP, AUX and SCONJ tags occur in the dictionaries.

Language	CHR11	GAR13	PLA16	AMB	FREQ	AMB+AE	FREQ+AE
am	0.3441	0.1392	0.1643 (0.00)	0.1595 (0.03)	0.2479 (0.00)	0.1651 (0.00)	0.2608 (0.01)
be	0.2366	0.3524	0.4234 (0.03)	0.4627 (0.03)	0.4805 (0.01)	0.4285 (0.02)	0.5027 (0.02)
br	0.3442	0.3119	0.3267 (0.01)	0.3449 (0.02)	0.3247 (0.01)	0.3375 (0.02)	0.3325 (0.01)
bxr	0.4432	0.5295	0.3140 (0.09)	0.3783 (0.07)	0.4114 (0.02)	0.4153 (0.09)	0.4372 (0.01)
fo	0.4048	0.5671	0.5928 (0.01)	0.5992 (0.00)	0.5559 (0.02)	0.6016 (0.00)	0.5341 (0.01)
hsb	0.1886	0.3657	0.3573 (0.06)	0.4306 (0.01)	0.3540 (0.01)	0.4125 (0.03)	0.3446 (0.01)
hy	0.3706	0.3821	0.4940 (0.02)	0.5131 (0.01)	0.4302 (0.02)	0.5061 (0.01)	0.4482 (0.01)
kk	0.451	0.4271	0.4801 (0.07)	0.4809 (0.06)	0.4524 (0.02)	0.5370 (0.02)	0.4469 (0.02)
kmr	0.3201	0.3501	0.3165 (0.06)	0.3898 (0.01)	0.3020 (0.01)	0.3865 (0.00)	0.2948 (0.01)
lt	0.383	0.4460	0.5251 (0.01)	0.5266 (0.02)	0.4813 (0.01)	0.5226 (0.01)	0.4857 (0.02)
mr	0.2522	0.3862	0.3670 (0.00)	0.3710 (0.01)	0.3781 (0.01)	0.3799 (0.01)	0.3808 (0.01)
mt	0.3126	0.3002	0.2208 (0.03)	0.2666 (0.04)	0.3326 (0.01)	0.2924 (0.06)	0.3544 (0.02)
ta	0.3275	0.2758	0.3302 (0.04)	0.3163 (0.05)	0.3193 (0.01)	0.3562 (0.00)	0.3259 (0.01)
te	0.5035	0.5062	0.4430 (0.06)	0.4746 (0.01)	0.5734 (0.02)	0.4888 (0.01)	0.5615 (0.01)
tl	0.2774	0.5274	0.4931 (0.01)	0.4924 (0.01)	0.5157 (0.02)	0.5075 (0.04)	0.4651 (0.04)
Average	0.3440	0.3911	0.3899	0.4138	0.4106	0.4225	0.4117

Table 2: POS tagging accuracy. For all neural models, the standard deviation is given in parentheses.

Language	CHR11	GAR13	AMB	FREQ
de -	0.29	0.42	0.47	0.44
es -	0.47	0.36	0.46	0.56
it -	0.39	0.32	0.41	0.48
pt -	0.49	0.34	0.37	0.53
sv -	0.32	0.42	0.43	0.49
Average -	0.39	0.37	0.43	0.50
de +	0.59	0.59	0.66	0.72
es +	0.63	0.74	0.72	0.80
it +	0.67	0.75	0.74	0.75
pt +	0.62	0.66	0.73	0.77
sv +	0.59	0.69	0.70	0.80
Average +	0.62	0.69	0.72	0.77

Table 3: High-resource POS tagging accuracy. “-”: comparable data to previous experiments; “+”: higher quality data.

been investigated exhaustively (Søgaard and Goldberg 2016; Augenstein and Søgaard 2018). An analysis of task combinations was performed by Bingel and Søgaard (2017). Ruder et al. (2017) presented a more flexible architecture, which learns what to share between the main and auxiliary tasks. Augenstein, Ruder, and Søgaard (2017) combined MTL with semi-supervised learning for strongly related tasks with different output spaces. However, work on combining sequence labeling main tasks and seq2seq auxiliary tasks is harder to find. Dai and Le (2015) pretrained an LSTM as part of a sequence autoencoder on unlabeled data to obtain better performance on a sequence classification task. However, they reported poor results for joint training. We obtain different results: an autoencoding seq2seq task is beneficial for low-resource POS tagging. Cross-lingual approaches have been used for a large variety of tasks, e.g., automatic speech recognition (Huang et al. 2013), entity recognition (Wang and Manning 2014), or parsing (Søgaard 2011; Naseem, Barzilay, and Globerson 2012; Ammar et al. 2016). In the realm of seq2seq models, work exists on cross-lingual machine translation (Dong et al. 2015; Johnson et al. 2017).

Another example is a character-based approach by Kann, Cotterell, and Schütze (2017) for morphological generation.

Conclusion

We analyzed state-of-the-art approaches for low-resource POS tagging of *truly* low-resource languages: POS tagging in these languages is still difficult because resources are limited and of poor quality, with average tagging accuracies well below 50%. Weakly supervised approaches only slightly outperform a state-of-the-art unsupervised baseline.

Acknowledgments

We would like to thank Barbara Plank, Isabelle Augenstein, and Johannes Bjerva for conversations on this topic. AS is funded by a Google Focused Research Award.

References

- Agić, Ž.; Hovy, D.; and Søgaard, A. 2015. If all you have is a bit of the Bible: Learning POS taggers for truly low-resource languages. In *ACL-IJCNLP*.
- Al-Rfou, R.; Perozzi, B.; and Skiena, S. 2013. Polyglot: Distributed Word Representations for Multilingual NLP. In *CoNLL*.
- Ammar, W.; Mulcaire, G.; Ballesteros, M.; Dyer, C.; and Smith, N. 2016. Many languages, one parser. *TACL* 4:431–444.
- Augenstein, I., and Søgaard, A. 2018. Multi-task learning of pairwise sequence classification tasks over disparate label spaces. In *NAACL*.
- Augenstein, I.; Ruder, S.; and Søgaard, A. 2017. Multi-task learning of keyphrase boundary detection. In *ACL*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Baldwin, T., and Lui, M. 2010. Language identification: The long and the short of the matter. In *NAACL*.
- Baldwin, T.; Pool, J.; and Colowick, S. M. 2010. PanLex and LEXTRACT: Translating all Words of all Languages of the World. In *COLING*.
- Ballesteros, M.; Dyer, C.; and Smith, N. A. 2015. Improved transition-based parsing by modeling characters instead of words with LSTMs. In *EMNLP*.

- Berg-Kirkpatrick, T.; Bouchard-Côté, A.; DeNero, J.; and Klein, D. 2010. Painless unsupervised learning with features. In *ACL*.
- Bingel, J., and Søgaard, A. 2017. Identifying beneficial task relations for multi-task learning in deep neural networks. In *EACL*.
- Bohnet, B.; McDonald, R.; Simões, G.; Andor, D.; Pitler, E.; and Maynez, J. 2018. Morphosyntactic tagging with a meta-biLSTM model over context sensitive token encodings. In *ACL*.
- Brown, P. F.; Desouza, P. V.; Mercer, R. L.; Pietra, V. J. D.; and Lai, J. C. 1992. Class-based n-gram models of natural language. *Computational linguistics* 18(4).
- Christodoulopoulos, C.; Goldwater, S.; and Steedman, M. 2010. Two Decades of Unsupervised POS Induction: How Far Have We Come? In *EMNLP*.
- Christodoulopoulos, C.; Goldwater, S.; and Steedman, M. 2011. A Bayesian Mixture Model for PoS Induction Using Multiple Features. In *EMNLP*.
- Currey, A., and Heafield, K. 2019. Incorporating source syntax into transformer-based neural machine translation. In *WMT*.
- Dai, A. M., and Le, Q. V. 2015. Semi-supervised sequence learning. In *NIPS*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
- Dong, D.; Wu, H.; He, W.; Yu, D.; and Wang, H. 2015. Multi-task learning for multiple language translation. In *ACL-IJCNLP*.
- Fang, M., and Cohn, T. 2017. Model Transfer for Tagging Low-resource Languages using a Bilingual Dictionary. In *ACL*.
- Faruqui, M., and Dyer, C. 2014. Improving Vector Space Word Representations Using Multilingual Correlation. In *EACL*.
- Garrette, D., and Baldridge, J. 2013. Learning a Part-of-Speech Tagger from Two Hours of Annotation. In *NAACL*.
- Goldwater, S., and Griffiths, T. 2007. A fully bayesian approach to unsupervised part-of-speech tagging. In *ACL*.
- Gouws, S., and Søgaard, A. 2015. Simple task-specific bilingual word embeddings. In *NAACL*.
- Hajič, J., and Zeman, D. 2017. Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies. In *CoNLL*.
- Hovy, D.; Plank, B.; Alonso, H. M.; and Søgaard, A. 2015. Mining for unambiguous instances to adapt POS taggers to new domains. In *NAACL*.
- Huang, J.-T.; Li, J.; Yu, D.; Deng, L.; and Gong, n. 2013. Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers. In *IEEE*.
- Johnson, M.; Schuster, M.; Le, Q.; Krikun, M.; Wu, Y.; Chen, Z.; Thorat, N.; Viégas, F.; Wattenberg, M.; Corrado, G.; Hughes, M.; and Dean, J. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *TACL* 5:339–351.
- Kann, K.; Bjerva, J.; Augenstein, I.; Plank, B.; and Søgaard, A. 2018. Character-level supervision for low-resource POS tagging. In *DeepLo*.
- Kann, K.; Cotterell, R.; and Schütze, H. 2017. One-shot neural cross-lingual transfer for paradigm completion. In *ACL*.
- Kirov, C.; Sylak-Glassman, J.; Que, R.; and Yarowsky, D. 2016. Very-large Scale Parsing and Normalization of Wiktionary Morphological Paradigms. In *LREC*.
- Li, S.; Graça, J. V.; and Taskar, B. 2012. Wiki-ly Supervised Part-of-Speech Tagging. In *EMNLP*.
- Meriardo, B. 1994. Tagging english text with a probabilistic model. *Computational Linguistics* 20(2):155–171.
- Moore, R. 2015. An improved tag dictionary for faster part-of-speech tagging. In *EMNLP*.
- Nadejde, M.; Reddy, S.; Sennrich, R.; Dwojak, T.; Junczys-Dowmunt, M.; Koeh, P.; and Birch, A. 2017. Predicting target language ccg supertags improves neural machine translation. In *WMT*.
- Naseem, T.; Barzilay, R.; and Globerson, A. 2012. Selective sharing for multilingual dependency parsing. In *ACL*.
- Nivre, J.; Agić, Ž.; Ahrenberg, L.; Aranzabe, M. J.; Asahara, M.; Atutxa, A.; Ballesteros, M.; Bauer, J.; et al. 2017. Universal dependencies 2.0. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Charles University.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep Contextualized Word Representations. In *ACL*.
- Petrov, S.; Das, D.; and McDonald, R. 2012. A Universal Part-of-Speech Tagset. In *LREC*.
- Plank, B., and Agić, Ž. 2018. Distant supervision from disparate sources for low-resource part-of-speech tagging. In *EMNLP*.
- Plank, B.; Søgaard, A.; and Goldberg, Y. 2016. Multilingual Part-of-Speech Tagging with Bidirectional Long Short-Term Memory Models and Auxiliary Loss. In *ACL*.
- Rei, M. 2017. Semi-supervised multitask learning for sequence labeling. In *ACL*.
- Ruder, S.; Bingel, J.; Augenstein, I.; and Søgaard, A. 2017. Sluice networks: Learning what to share between loosely related tasks. *arXiv:1705.08142*.
- Sennrich, R., and Haddow, B. 2016. Linguistic input features improve neural machine translation. In *WMT*.
- Søgaard, A., and Goldberg, Y. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *ACL*.
- Søgaard, A. 2011. Data point selection for cross-language adaptation of dependency parsers. In *ACL-HLT*.
- Stratos, K.; Collins, M.; and Hsu, D. 2016. Unsupervised Part-Of-Speech Tagging with Anchor Hidden Markov Models. *TACL* 4.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Täckström, O.; Das, D.; Petrov, S.; McDonald, R.; and Nivre, J. 2013. Token and Type Constraints for Cross-Lingual Part-of-Speech Tagging. *TACL* 1:1–12.
- Vosoughi, S.; Vijayaraghavan, P.; and Roy, D. 2016. Tweet2Vec: Learning tweet embeddings using character-level CNN-LSTM encoder-decoder. In *SIGIR*.
- Wang, M., and Manning, C. D. 2014. Cross-lingual pseudo-projected expectation regularization for weakly supervised learning. *TACL* 2:55–66.
- Wang, P.; Qian, Y.; Soong, F. K.; He, L.; and Zhao, H. 2015. A unified tagging solution: Bidirectional LSTM recurrent neural network with word embedding. *arXiv:1511.00215*.
- Wisniewski, G.; Pécheux, N.; Gahbiche-Braham, S.; and Yvon, F. 2014. Cross-Lingual Part-of-Speech Tagging through Ambiguous Learning. In *EMNLP*.
- Zeman, D.; Hajič, J.; Popel, M.; Potthast, M.; Straka, M.; Ginter, F.; Nivre, J.; and Petrov, S. 2018. CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies. In *CoNLL*.