

Morphologically Aware Word-Level Translation

Paula Czarowska¹ Sebastian Ruder³ Ryan Cotterell^{1,2} Ann Copestake¹

¹University of Cambridge ³DeepMind ²ETH Zürich

pjc211@cam.ac.uk sebastian@ruder.io

ryan.cotterell@inf.ethz.ch aac10@cam.ac.uk

Abstract

We propose a novel morphologically aware probability model for bilingual lexicon induction, which jointly models lexeme translation and inflectional morphology in a structured way. Our model exploits the basic linguistic intuition that the lexeme is the key lexical unit of meaning, while inflectional morphology provides additional syntactic information. This approach leads to substantial performance improvements—19% average improvement in accuracy across 6 language pairs over the state of the art in the supervised setting and 16% in the weakly supervised setting. As another contribution, we highlight issues associated with modern BLI that stem from ignoring inflectional morphology, and propose three suggestions for improving the task.

1 Introduction

The ability to generalize to rare and unseen morphological variants of a known word lies at the heart of translation. For instance, a capable human Spanish–English translator would find translating the exceptionally rare form *tosed* (2nd person plural imperative form of ‘to cough’) as straightforward as translating the infinitive *toser*—despite the fact that *tosed* is so infrequent that many native Spanish speakers may have never encountered the form themselves. Given how basic this generalization ability is for humans, one should expect a good bilingual lexicon inducer to exhibit a similar capacity to generalize. In other words, the model should translate infrequent, regular forms of a lexeme as accurately as it translates a lexeme’s most common forms.

Nevertheless, current approaches to bilingual lexicon induction (BLI) fall short of this desideratum. In a recent study, Czarowska et al. (2019) reveal that the performance of state-of-the-art bilingual lexicon inducers degrades severely when translating less frequent inflected forms—even for the most common lexemes. The problem is severe: In the case of inducing a French–Spanish bilingual lexicon, the model of Ruder et al. (2018) correctly translates infinitives 50.6% of the time, but correctly translates the 2nd person plural imperative forms only 1.5% of the time. Motivated by this disparity, this work introduces a novel morphologically aware probability model for BLI that jointly models lexeme translation and inflectional morphology in a structured way. Our model exploits the basic linguistic intuition that the lexeme is the core lexical unit of meaning, while inflectional morphology provides additional syntactic information on top of it (Haspelmath and Sims, 2013). It follows that we should ignore this syntactic information when translating at the word level and handle morphological inflection with a different component of the model.

The empirical portion of our paper describes experiments on French, Italian and Spanish. We find our joint model substantially improves over several strong baselines on the BLI task. When evaluating on held-out lexemes,¹ we observe an average performance improvement of 19% and 16% over the previous state of the art (Artetxe et al., 2018b) in the supervised and weakly supervised settings, respectively. In addition, we propose a simple heuristic to further boost performance: Inspired by the dual-route hypothesis (Pinker and Prince, 1994; Pinker, 1998) and the network model of morphological processing (Bybee, 1985; Bybee, 1995b), we translate high-frequency forms, which are most likely to exhibit irregularity, directly (without going through the lexeme) and reserve our morphologically aware model for low-frequency forms. This heuristic gives us a further 2% improvement.

¹This is the case when we have never seen any inflected form of a given lexeme in the training data.

2 Background

2.1 Morphological Inflection

Morphological inflection is the systematic alteration of the word form that adds specific morpho-syntactic information to the lexeme, e.g. tense, case and number. English is weakly inflected with only a few forms per lexeme, and, in that respect, it differs from many other languages. A higher degree of inflection is indeed the norm among the world’s languages (Dryer and Haspelmath, 2013). In this work, we distinguish the terms lexemes, inflected forms and lemmata. A **lexeme** is an abstract concept that represents the core meaning shared by a set of inflected forms. An **inflected form** is an individual morphological variant that belongs to a given lexeme. As an example, the lexeme RUN has the inflected forms *run*, *runs*, *ran*, and *running*. The **lemma**, also called the citation form, is an inflected form that lexicographers have chosen to be representative of the lexeme. For example, the lexeme RUN’s lemma is *run*. In many languages, the infinitive is the verbal lemma and the nominative singular is the nominal lemma. We consider a **lexicon** of a language to be a set of inflected forms.²

2.2 Bilingual Lexicon Induction

In the NLP literature, the BLI task is to translate a given list of source-side word forms into the most appropriate corresponding target-side word forms. It dates back to 1990s. The first data-driven experiments on parallel corpora made use of word-alignment techniques (Brown et al., 1990; Kupiec, 1993; Smadja and McKeown, 1994). Such approaches were later extended to operate on non-parallel or even unrelated texts by leveraging the correlation between word co-occurrence patterns in different languages (Rapp, 1995; Fung and Lo, 1998; Fung, 1998; Koehn and Knight, 2002). Apart from the distributional signal, the early approaches make use of other monolingual clues, e.g. word spelling, cognates or word frequency.³ More recent approaches leverage the distributional signal in word embeddings without any explicit linguistic clues. Many current models (Mikolov et al., 2013; Ruder et al., 2019) learn a linear transformation between two monolingual word embedding spaces, often guided by an initial set of seed translations. This seed dictionary frequently spans several thousand word pairs (Mikolov et al., 2013; Xing et al., 2015; Lazaridou et al., 2015; Artetxe et al., 2016) but one can also provide weaker supervision, through listing only identical strings or shared numerals (Artetxe et al., 2017; Søgaard et al., 2018). For unsupervised BLI, the initial translations may also be induced automatically through exploiting the structure of the monolingual embedding spaces (Zhang et al., 2017; Conneau et al., 2018; Artetxe et al., 2018b). We focus on supervised and weakly supervised BLI which outperform unsupervised approaches (Glavaš et al., 2019). The BLI models are typically evaluated using the precision@*k* metric, which tells us how many times the correct translation of a source form is among the *k*-best candidates returned by the model. In this work we exclusively consider the precision@1 metric, which is the least forgiving.

2.3 Morphological Inflection: A Challenge for BLI

Most datasets for BLI operate at the level of inflected forms and impose no restriction on the morpho-syntactic category of translated words. From a lexicographer’s standpoint, this choice is unusual. Dictionaries generally only list lemmata; inflected forms are rarely listed. Thus, BLI is *not* the task of inducing a dictionary in the strict sense of the word. The authors have found this a common misconception in the literature and among NLP researchers.

Despite the assumption that inflectional morphology is present in the lexicon, most BLI datasets list only a handful of inflected forms per lexeme. This is due to frequency restrictions imposed when the datasets are created and the consequences can be quite severe; Czarnowska et al. (2019) reveal that the MUSE test dictionaries (Conneau et al., 2018) for Romance language pairs cover, on average, only 3% of paradigms for verbs they contain. In contrast, we assert that BLI models should be trained and

²This full-listing view of the lexicon is taken, among others, by Jackendoff (1975), McClelland et al. (1987) and Bybee (1995b), and stands in opposition to alternative views of the lexicon being comprised of only the unpredictable items (e.g. Prasada and Pinker (1993) Pinker and Prince (1994)).

³For more references we refer the reader to the survey of Ruder et al. (2019).

evaluated on datasets that contain a more representative range of morphological inflections. We use the term **morphologically enriched dictionary** for such bilingual lexicons (see §5.1).

To our knowledge, we are the first to explicitly model inflectional morphology in BLI. Closest to our endeavor, Yang et al. (2019) address morphology in BLI by incorporating grammatical information learned by a pre-trained denoising language model, while Riley and Gildea (2018) enhance the projection-based approach of Artetxe et al. (2017) with orthographic features to improve performance on BLI for related languages.

3 A Joint Model for Morphologically Aware Word-level Translation

The primary contribution of this work is a morphologically aware probabilistic model for word-level translation. Our model exploits a simple intuition: Because the core unit of meaning is the lexeme, one should translate *through* the lexeme and then inflect the word according to the target language’s morphology.

Notation. In the task of BLI, we consider a source language s and a target language t . The goal is to translate inflected forms in a source language $\iota_s \in L_s$ to inflected forms in a target language $\iota_t \in L_t$, where L_s and L_t are the source-side and the target-side lexicons, respectively. We denote source-side lemmata as $\lambda_s \in L_s$ and target-side lemmata as $\lambda_t \in L_t$. We use $\tau_s \in T_s$ for the morpho-syntactic tag of the source form ι_s and $\tau_t \in T_t$ for the tag of the target form ι_t , where T_s and T_t are the sets of possible tags in the source and the target language. Finally, let $e : L \rightarrow \mathbb{R}^N$ be a function which takes as input a word and returns its pre-trained monolingual word embedding.

Model. We construct a joint probability model for morphologically aware word-level translation. Using the notation defined in the previous paragraph, we formally define the model as

$$p(\boldsymbol{\pi}_t, \lambda_t, \tau_s, \lambda_s, | \iota_s) = \underbrace{\left[\prod_{\langle \iota_t, \tau_t \rangle \in \boldsymbol{\pi}_t} p(\iota_t, \tau_t | \lambda_t, \tau_s) \right]}_{\text{synthesizer}} \underbrace{p(\lambda_t | \lambda_s)}_{\text{translator}} \underbrace{p(\tau_s, \lambda_s | \iota_s)}_{\text{analyzer}} \quad (1)$$

where $\boldsymbol{\pi}_t$ is a *set* of valid target-language translations. The joint distribution is factorized into three parts: a synthesizer, a translator and an analyzer. In the next three sections, we define each of these distributions. Note that the model, as defined in eq. (1), can provide a *set* of valid translations $\boldsymbol{\pi}_t$. For languages with very similar morphological systems, this set will often have one element (with the same morpho-syntactic description), while for more distinct languages it will contain a number of inflected forms.

3.1 The Synthesizer: $\prod_{\langle \iota_t, \tau_t \rangle \in \boldsymbol{\pi}_t} p(\iota_t, \tau_t | \lambda_t, \tau_s)$

The synthesizer produces a set of valid target-side inflected forms and their tags $\boldsymbol{\pi}_t$ given a target-side lemma λ_t and a source-side tag τ_s . We formally define this distribution as follows:

$$\prod_{\langle \iota_t, \tau_t \rangle \in \boldsymbol{\pi}_t} p(\iota_t, \tau_t | \lambda_t, \tau_s) = \prod_{\langle \iota_t, \tau_t \rangle \in \boldsymbol{\pi}_t} \underbrace{p(\iota_t | \tau_t, \lambda_t)}_{\text{inflector}} \underbrace{p(\tau_t | \tau_s)}_{\text{tag translator}} \quad (2)$$

The joint distribution over forms and tags is factored into two parts. The first part, the inflector, produces an inflected form ι_t given a lemma λ_t and a morphological tag τ_t . This problem has been well studied in the NLP literature (Cotterell et al., 2016; Cotterell et al., 2017). The second part, tag translator, determines the possible target-side morphological tags that are compatible with the features present in the source tag.

In principle, our model is compatible with any probabilistic inflector. In this paper, we employ the model of Wu et al. (2019b), which obtained the single-model state of the art at the time of experimentation (McCarthy et al., 2019). The model has a latent character-level monotonic alignment between the source and target inflections that is jointly learned with the transducer and is, in effect, a neuralized version of a hidden Markov model for translation (Vogel et al., 1996).

In this work we focus on closely related languages and make a simplifying assumption that there exists a single most-plausible translation for each inflected form.⁴ We formalize the tag translator using an indicator function:

$$p(\tau_t | \tau_s) = \begin{cases} 1 & \text{if } \tau_t = \tau_s \\ 0 & \text{if } \tau_t \neq \tau_s \end{cases}$$

For experiments with more distant language pairs one can define $p(\tau_t | \tau_s)$ to be a multi-label classifier.

3.2 The Translator: $p(\lambda_t | \lambda_s)$

As our translator, we construct a log-bilinear model that yields a distribution over all elements in the target lexicon. We assume the existence of both source- and target-side embeddings. For notational simplicity, we use \mathbf{e} to define the embedding function for both the source and the target language, although in practice these look-up functions are distinct. The model has a single matrix of parameters: $\Omega \in \mathbb{R}^{N_t \times N_s}$ where N_s is the source embedding dimensionality and N_t the target embedding dimensionality. Our translator is defined as the following conditional model

$$p(\lambda_t | \lambda_s) = \frac{1}{Z(\lambda_s)} \exp(\mathbf{e}(\lambda_t)^\top \Omega \mathbf{e}(\lambda_s)) \quad (3)$$

where the normalizer is defined as

$$Z(\lambda_s) = \sum_{\lambda'_t \in L_t} \exp(\mathbf{e}(\lambda'_t)^\top \Omega \mathbf{e}(\lambda_s)) \quad (4)$$

Note that this log-bilinear model differs from most embedding-based bilingual lexicon inducers which predict embeddings, rather than words. For example, Ruder et al. (2018)’s approach contains a multivariate Gaussian over the target-language’s embedding space.⁵

Orthogonal Regularization. During training we employ a special regularization term on the parameter matrix Ω . Specifically, we use

$$R(\Omega) = \alpha \left\| \Omega^\top \Omega - \mathbf{I} \right\|_F \quad (5)$$

with a tunable “strength” hyperparameter $\alpha \in \mathbb{R}_{\geq 0}$. This term encourages the translation matrix to be orthogonal, which has led to consistent gains in past work (Xing et al., 2015; Artetxe et al., 2016; Ruder et al., 2018).

3.3 The Analyzer: $p(\lambda_s, \tau_s | \iota_s)$

For our probabilistic analyzer we use the same hard attention model as in the inflector. The model predicts both the lemma and the morphological tag; the output is a morphological tag followed by a special end-of-tag character and a sequence of lemma characters.

4 A Frequency-Based Heuristic

As defined in §3, the model handles regular and irregular word forms in the same manner. This can be problematic, as the analysis and synthesis modules have no means of handling irregular morphology, beyond the irregular forms they have been exposed to during training. Guided by this insight, we propose an alternative version of our model, which employs a special treatment for forms likely to have irregular morphology—the most frequent forms (Bybee, 1995a; Baayen and Lieber, 1996; Wu et al., 2019a). We term this extension the **hybrid model**. It employs a frequency-based heuristic and translates the source form through its lemma *only* if the lemma is more frequent.⁶ Otherwise, it translates the inflected form

⁴Indeed, the dictionaries of Czarnowska et al. (2019) on which we experiment also make this assumption.

⁵Note that the translator’s distribution, as defined in eq. (3) and eq. (4), is over all inflected forms in the target lexicon. An alternative would be to define a distribution over lemmata only. However, this would require filtering out all non-lemma forms from the target embedding matrix, which is not trivial. In our preliminary experiments, we observed that this can lead to a further performance increase.

⁶We rely on the order of FASTTEXT embeddings for the relative ranking of inflected forms.

directly, using only the translator component. For example, the Spanish form *pide*—an irregular inflected form of *pedir* (‘to ask for’)—would be translated directly, since *pedir* is less frequent.

In a broader context of morphological processing, different handling of a form depending on its frequency or regularity can be linked to the dual-route hypothesis (Pinker and Prince, 1994; Pinker, 1998), which posits that regular and irregular inflection are handled by different cognitive mechanisms, or the works of Baayen (1992), Baayen (1993) and Hay (2001), which have loosely inspired this heuristic.⁷

5 Experimental Setup

Our evaluation involves 3 Romance languages which exhibit a higher degree of inflection and are commonly experimented on within BLI—French, Spanish and Italian. Because these languages come from the same branch of the Indo-European family, the results serve as an empirical upper-bound on BLI.

5.1 BLI Datasets

Czarnowska et al. (2019). As discussed in §2.3, given that inflectional morphology is present in the induced lexicon, BLI models should be trained and evaluated on datasets which list a range of compatible inflected form pairs for every source-target lexeme pair. At this time, the dictionaries of Czarnowska et al. (2019) are the only publicly available resource that meets this criterion, and, for this reason, they are the most important evaluation benchmark used in this work. The dictionaries were generated based on Open Multilingual WordNet (Bond and Paik, 2012), Extended Open Multilingual WordNet (Bond and Foster, 2013) and UniMorph⁸ (Kirov et al., 2016; McCarthy et al., 2020), a resource comprised of inflectional word paradigms for 107 languages. The dictionaries only list parts of speech that undergo inflection in either the source or the target language; these are nouns, adjectives and verbs in the Romance languages.

Conneau et al. (2018). MUSE (Conneau et al., 2018) was generated using an “internal translation tool” and is one of the few other resources which covers pairs of Romance languages. However, it is skewed towards most frequent forms: The vast majority of forms in MUSE are ranked in the top 10k of the vocabularies in their respective languages, causing it to omit many morphological variants of words. The dataset also suffers from other issues, such as a high level of noise coming from proper nouns (Kementchedjheva et al., 2019). Thus, we do not view this resource as a reasonable benchmark for BLI.

5.2 Baselines

Artetxe et al. (2016). They learn an orthogonal linear transformation matrix between the source language space and the target language space, after length-normalizing and mean-centering the monolingual embedding matrices. Their method is fully supervised and works best with large amounts of training data (several thousand translation pairs).

Ruder et al. (2018). They introduce a weakly supervised, self-learning model, which can induce a dictionary given only a handful of initial, seed translations. This is achieved by iteratively alternating between two steps: an inflected form alignment step and a mapping step. The first is comprised of finding a matching in a bipartite weighted graph in which source forms constitute one set and target forms the other set. During the mapping step the resulting alignment is used to learn a better projection from the source language space into the target space—this is done by solving the orthogonal Procrustes problem.⁹

Artetxe et al. (2018b). They propose a fully unsupervised approach to BLI. The starting point for their model is an automatic initialization of a seed dictionary which exploits the structural similarity of the monolingual embeddings. Like Ruder et al. (2018), their model also utilizes self-learning, but learns mappings in both directions and employs a range of additional training techniques. Despite being unsupervised, this model constitutes the state of the art on many BLI datasets (e.g. Dinu et al. (2015)’s

⁷In particular, Hay (2001) advocates the importance of *relative* (vs. absolute) frequency for lexical access, arguing that it is more beneficial to define a frequent complex form as one that is more frequent than its base, regardless of how often it is used.

⁸<https://unimorph.github.io/>

⁹Given two matrices \mathbf{A} and \mathbf{B} , the orthogonal Procrustes problem is to find an orthogonal matrix \mathbf{Q} which most closely maps from \mathbf{A} to \mathbf{B} : $\mathbf{Q} = \arg \max_{\Omega} \|\Omega \mathbf{A} - \mathbf{B}\|_F$ subject to $\Omega^T \Omega = \mathbf{I}$

and Artetxe et al. (2018a)’s), outperforming even the supervised approaches. Thus, we include it in our evaluation and compare it to supervised and semi-supervised approaches.

5.3 Skyline

We also consider a version of our model which uses an oracle analyzer—the source lemma λ_s and tag τ_s are known *a priori*. The skyline provides an upper-bound of performance—to wit, what performance would be achievable if the model had had access to more information about the translated source form.

5.4 Experimental Details

We implemented all models in PyTorch (Paszke et al., 2019), adapting the code of Wu et al. (2019b) for the transducers (analyzer and inflector). Throughout our experiments we used the Wikipedia FASTTEXT embeddings (Grave et al., 2018), which we length-normalized and mean-centered before training the models. As is standard, we trained all translators on the top 200k most frequent word forms in the vocabularies of both languages. To evaluate on very rare forms present in the dictionaries of Czarnowska et al. (2019) which are out-of-vocabulary (OOV) for FASTTEXT, we created an OOV FASTTEXT embedding for every OOV form that appears in a union of WordNet and UniMorph and appended those representations to the original embedding matrices.¹⁰ We evaluated all models using precision@1 as a metric, which is equivalent to accuracy. At evaluation, for all models we used cosine as a measure of similarity between two word embeddings.¹¹

Estimating the Model. We estimate the parameters of our models to maximize the log-likelihood of the training data. In the supervised case, we are able to estimate the parameters of the different components of the model independently. For every language pair we trained a separate translator, on the initial seed dictionary, as well as a separate analyzer and inflector on UniMorph entries—in the source language for the analyzers and the target language for the inflectors. Importantly, to ensure that the transducers are never trained on forms they will see at evaluation, we excluded the entries present in the test or development split of the considered evaluation dataset.¹²

For inflection, we make the assumption that there always exists one unique inflected form for a given lemma and tag. However, in the case of analysis, due to syncretism, there often exists a number of plausible interpretations. As a result, in the training data there might be a number of correct lemma–tag combinations for every input form. At training, we select only one of those possible analyses. We found that this approach works better than training the model on all possible analyses as targets. As a consequence, the analyzers might be biased towards specific morpho-syntactic interpretations of syncretic forms, which, down the line, may hurt the performance of our model. Indeed, we view this as a trade-off between having a more accurate but biased analyzer that limits the possible translations produced by the model, and a more noisy analyzer which can result in more varied translations.

Hyperparameters. The models were trained with Adam (Kingma and Ba, 2014), using an initial learning rate of 0.001 for the transducers, and 0.05 for most translators (see §A for a detailed breakdown). We halved the learning rate after every epoch for which the development loss increased and utilized early stopping (with the min. learning rate of 1×10^{-8}). For the transducers we also applied gradient clipping with a maximum gradient norm of 5. For Ruder et al. (2018), we set the number of candidate target words considered for each source word during matching to 15 and constrained the matching to consider only 40k most frequent forms.

¹⁰The FASTTEXT framework allows for creating a vector of an unseen form by summing the vectors of all character n -grams present in that form.

¹¹CSLS (Conneau et al., 2018) is an alternative retrieval method and can lead to better results.

¹²Since the MUSE dataset contains word forms associated with parts of speech not present in UniMorph, i.e. adverbs, determiners and pronouns; in our experiments on those datasets we additionally enhanced the transducers’ training data with forms associated with those POS. All these additional forms are base forms and the models are trained to analyze them as such and not to inflect them.

Model	FRA-ITA		FRA-SPA		ITA-FRA		ITA-SPA		SPA-FRA		SPA-ITA	
	VOC	ALL	VOC	ALL	VOC	ALL	VOC	ALL	VOC	ALL	VOC	ALL
Experiments on the morphologically enriched dictionaries (Czarnowska et al., 2019)												
Train/test split sizes	94k/21k		85k/22k		96k/31k		70k/31k		86k/26k		70k/24k	
Supervised setting: full train dictionary splits were used for training.												
Artetxe et al. (2016)	49.4	28.3	49.1	29.3	40.8	21.4	43.6	22.6	46.9	25.8	47.0	27.8
Ruder et al. (2018)	48.3	26.8	47.8	26.7	40.5	21.0	42.7	21.1	45.5	24.2	46.6	27.1
Our base model	50.5	47.8	48.6	45.1	42.8	35.2	46.0	38.6	51.5	44.3	55.4	51.6
Our hybrid model	56.7	50.6	54.6	47.9	47.2	36.9	50.3	40.2	55.2	45.9	58.3	52.9
Oracle	58.1	52.2	56.3	49.3	48.7	37.3	51.8	41.0	56.7	46.8	61.2	54.3
Weakly supervised setting: pairs of forms identical in both languages were used as a seed dictionary.												
Artetxe et al. (2016)	44.7	25.2	45.2	24.8	35.8	18.5	40.8	19.5	42.2	22.2	44.1	24.9
Ruder et al. (2018)	48.4	26.9	47.9	26.6	40.7	20.9	42.7	21.0	45.6	24.2	46.4	26.9
Our base model	46.7	43.2	45.8	42.0	41.1	33.8	44.7	36.4	49.2	40.8	52.9	47.3
Our hybrid model	52.2	45.7	51.2	44.3	44.7	35.2	48.9	37.9	52.5	42.3	55.8	48.7
Oracle	53.6	47.0	53.0	45.6	46.9	36.0	50.3	38.6	54.3	43.6	58.3	50.1
Unsupervised setting: no seed dictionary was provided.												
Artetxe et al. (2018b)	49.4	28.2	49.2	27.4	38.5	20.0	40.6	20.0	43.2	23.0	47.4	28.2
Experiments on the MUSE (Conneau et al., 2018) (Train/test split sizes: 5000/1500)												
Supervised setting: full train dictionary splits were used for training.												
Artetxe et al. (2016)	78.9	78.9	78.9	78.9	83.5	83.5	84.6	84.6	82.1	82.1	79.5	79.4
Ruder et al. (2018)	76.4	76.4	76.6	76.6	83.9	83.9	83.7	83.7	82.5	82.5	78.9	78.9
Our base model	58.2	58.2	50.7	50.7	63.1	63.1	61.7	61.7	58.7	58.7	56.1	56.1
Our hybrid model	75.1	75.1	71.6	71.6	79.7	79.7	81.4	81.4	75.1	75.1	75.4	75.3
Unsupervised setting: no seed dictionary was provided.												
Artetxe et al. (2018b)	73.1	73.1	74.0	74.0	77.7	77.7	81.0	81.0	77.7	77.7	76.6	76.6

Table 1: BLI results: VOC considers in-vocabulary forms and ALL considers the full dictionary, including OOVs.

Decoding. We decode the model with greedy search, i.e. beam search with a beam size of 1. In all experiments we return the single most suitable form.

Data Requirements. As for other projection-based BLI approaches, the translator component needs to be trained on monolingual embeddings and an initial seed dictionary, which can be generated automatically. The *only* additional resource required to train the full model is UniMorph, or a similar morphological database, which is used to train the transducers. Although this extra requirement could be a limitation in some cases, such morphological lexicons are available in an increasingly large number of languages because they are a by-product of descriptive linguists’ efforts to document the world’s languages.

6 Results

We consider the fully supervised and the weakly supervised setting. Weak supervision in the case of BLI refers to populating the seed dictionary with identically spelled strings (Søgaard et al., 2018). In Table 1, we present our experimental results on the whole evaluation dictionary, including the words which are out-of-vocabulary for FASTTEXT (ALL), as well as on the in-vocabulary forms only (VOC). When evaluated on the morphologically enriched resource, our proposed approach leads to substantial performance gains over the baseline models, on every language pair. In the supervised setting, for the in-vocabulary words we note an average 3% improvement in the case of our base model and a 8% improvement in the case of our hybrid model over the best performing baseline (Artetxe et al., 2016). In the full-vocabulary experiments, the improvements reach 18% for our base model and 20% for the hybrid. In the weakly supervised setting, the performance gains remain similarly high—for our base model they reach 2% and 16%, respectively, compared to the best baseline (Artetxe et al., 2018b), while for the hybrid model we observe a 6% and 18% improvement, respectively. Our experiments with an oracle analyzer




Figure 1: The models’ performance on the full dictionary of Czarnowska et al. (2019) (left hand side of the plots) and across different word frequency bins, e.g. 0–10k bin is for the top 10k inflected forms.

demonstrate that even larger gains are possible: the oracle consistently outperforms every other model across all evaluation conditions.

When evaluated on MUSE (Conneau et al., 2018), the hybrid model is competitive with the baselines, while our base model performs worse. This is not surprising given that MUSE contains only the most frequent inflected forms which, as discussed in §4 and as we demonstrate later in this section, are the weakness of our base model. The hybrid model can also suffer from misanalyzing frequent irregular forms of even more frequent lemmata, such as *puede* or *pudo*—both forms of Spanish *poder* (‘to be able to’).

Frequency Breakdown. In Figure 1 we present the performance breakdown across 10 frequency bins for our base model, its hybrid version, the oracle and the model of Artetxe et al. (2018b) for an example language pair: Italian–Spanish. The plot displays results for the supervised setting, for the dictionary of Czarnowska et al. (2019).¹³ As per the plot, our method proves to be particularly beneficial for translating forms of medium to low frequency; in contrast to the baseline model, for which the performance continuously drops as the forms become less frequent, the performance of our models drops initially, but then plateaus at around 40% accuracy (for the in-vocabulary forms). Figure 1 also exposes a weakness of our base model—it leads to substantial performance drop for the two highest frequency bins. Notably, the hybrid model does not suffer from this limitation. Indeed, the fact that the shape of the hybrid’s plot closely resembles that for the oracle suggests the heuristic we used successfully identifies irregular forms, which are hard to analyze.

Morphology Breakdown. In Table 2 we present the translation accuracy on French–Spanish, in the supervised setting, for the model of Artetxe et al. (2018b) and the hybrid model across a range of different morpho-syntactic categories; the models are only evaluated on source forms belonging to a particular paradigm slot. We observe that our approach leads to improvements on all but two verbal paradigm slots, while for the adjectives and nouns the performance is competitive with the baselines. For many categories the improvements are very substantial; e.g., for the in-vocabulary 2nd person plural imperative forms of a verb (V;POS;IMP;2;PL) and 2nd person plural present subjunctive forms (V;SBJV;PRS;2;PL) the accuracy improves by 55%.

7 A Critique of Modern BLI

The model we develop in §3 stems from a desire to better integrate inflectional morphology into current state-of-the-art models for BLI. On one hand, the empirical findings we discuss in §6 indicate this attempt

¹³MUSE dictionaries (Conneau et al., 2018) only contain most frequent forms. Thus, we do not show the performance breakdown for that resource.

Morphology	(Artetxe et al., 2018b)		Our hybrid model		Morphology	(Artetxe et al., 2018b)		Our hybrid model	
	VOCAB	ALL	VOCAB	ALL		VOCAB	ALL	VOCAB	ALL
ADJ;FEM;PL	62.54	61.38	58.42	57.50	V;PRS;3;PL	47.55	42.11	47.80	43.64
ADJ;FEM;SG	59.52	58.90	57.61	57.02	V;PRS;3;SG	45.75	43.76	46.21	44.20
ADJ;MASC;PL	60.99	58.91	60.46	58.23	V;PST;1;SG;IPFV	2.30	11.43	58.62	46.15
ADJ;MASC;SG	63.43	61.72	61.01	59.53	V;PST;1;SG;PFV	7.02	1.76	38.60	43.30
N;PL	55.21	51.81	54.62	51.44	V;PST;2;PL;IPFV	0.00	0.88	42.86	44.84
N;SG	57.85	56.30	55.26	53.79	V;PST;2;PL;PFV	0.00	0.00	100.00	46.81
V;NFIN	48.51	47.05	51.03	49.45	V;PST;3;PL;IPFV	60.89	47.59	59.41	48.03
V;POS;IMP;1;PL	38.13	18.02	64.03	43.96	V;PST;3;PL;PFV	48.80	37.72	58.80	47.59
V;POS;IMP;2;PL	1.54	2.42	56.92	45.71	V;PST;3;SG;IPFV	57.78	48.80	55.39	47.70
V;PRS;1;SG	43.14	39.78	43.89	41.54	V;PST;3;SG;PFV	38.27	32.39	51.96	46.17
V;PRS;1;SG	43.14	39.78	43.89	41.54	V;SBJV;PRS;1;PL	20.51	4.18	66.67	45.93
V;PRS;2;PL	0.00	2.42	59.34	45.71	V;SBJV;PRS;1;SG	43.52	39.12	47.93	44.18
V;PRS;2;SG 2	37.05	25.71	48.66	40.66	V;SBJV;PRS;2;PL	0.00	0.88	55.88	44.62

Table 2: Performance on translating source forms belonging to different morpho-syntactic categories for French–Spanish.

was a success, but, on the other, our more nuanced conclusion is that the task of BLI, as currently researched in NLP, is ill-defined with respect to inflectional morphology. Indeed, the authors suggest that BLI needs redirection going forward.

The recent trend in BLI research is to remain data-driven and to avoid specialist linguistic annotation. Current projection-based approaches to BLI depend heavily on the assumption that the lexicons of different languages are approximately isomorphic (Mikolov et al., 2013; Miceli Barone, 2016). However, given the immense variation in morphological systems of worlds’ languages, this assumption is *prima facie* false. Consider the simple contradiction of Spanish and English, where the first exhibits much more morphological inflection than the latter; there can be no one-to-one alignment between the words in those two lexicons. The failure of the isomorphism assumption has been discussed and addressed in many recent works on cross-lingual word embeddings (Søgaard et al., 2018; Nakashole and Flauger, 2018; Ormazabal et al., 2019; Vulić et al., 2019; Patra et al., 2019). However, none of those studies directly target inflectional morphology. In this work we highlight that inflectional morphology complicates BLI and NLP researchers should strive to develop a cleaner way to integrate it into their models. We contend the models we present make progress in this direction but there is still a long way to go. We now make three concrete suggestions for BLI going forward. The first two involve engaging with morphology more seriously and are extensions to the ideas in this paper. The third focuses on backing away from morphology.

More Fine-Grained Lexicons. Our first proposal is to create more elaborate morphological dictionaries, in the style of those by Czarnowska et al. (2019). The primary draw-back of this suggestion is that such resources are tedious to create. Czarnowska et al. (2019) focus on genetically related languages, for which inflectional morphological systems are easily compatible. However, this is often not the case. Every morphological system carves up the semantic space in its own way, e.g. there is no good German equivalent of Polish verbal aspect. Thus, such elaborate resources should be carefully crafted and specify lemmata, inflected forms and tags for both the source and the target language. It follows, modulo polysemy and other lexical ambiguity, that the task would be well defined.

Contextual Word-level Translation. Another suggestion for future work is to contextualize word-level translation. Indeed, translation of syntactic features without context is somewhat unusual—e.g. in some contexts a feminine adjective in Spanish might not be translated as feminine in Italian, because a feminine noun in Spanish might be masculine in Italian. Arguably, much of the morphological ambiguity (and lexical ambiguity) present in modern BLI can be resolved by a word’s source-side context. However, identifying the amount and type of context sufficient to disambiguate possible translations is non-trivial. An additional point is that, in this scenario, BLI starts to approach full machine translation.

Lexeme-Level Translation. The final suggestion is for the task of BLI to ignore morphology. This would mean filtering the training and test lexicons to ensure that only lemmata exist. Such filtering can be

performed at the type level with a list of valid lemmata. This third suggestion would entail a return of BLI to the spirit of the task: the induction of a bilingual dictionary. However, creating a lemma-only resource requires additional language-specific knowledge.

8 Conclusion

We propose a novel model for bilingual lexicon induction which jointly models lexeme translation and inflectional morphology in a structured way. Our model improves handling of less frequent, morphologically complex forms, especially if they belong to large inflectional paradigms. In our experiments on morphologically enriched dictionaries, we observe substantial performance improvements over the state of the art, particularly prominent for rare verbal paradigm slots.

9 Acknowledgments

We thank the anonymous reviewers for their thoughtful comments. We would also like to thank Chris Dyer for his helpful suggestions. We acknowledge that Paula is supported by the Cambridge Trust Vice-Chancellor's and Selwyn College scholarship.

References

- Mikel Artetxe, Gorika Labaka, and Eneko Agirre. 2016. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2289–2294.
- Mikel Artetxe, Gorika Labaka, and Eneko Agirre. 2017. Learning bilingual word embeddings with (almost) no bilingual data. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462. Association for Computational Linguistics.
- Mikel Artetxe, Gorika Labaka, and Eneko Agirre. 2018a. Generalizing and improving bilingual word embedding mappings with a multi-step framework of linear transformations. In *AAAI Conference on Artificial Intelligence*.
- Mikel Artetxe, Gorika Labaka, and Eneko Agirre. 2018b. 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. Association for Computational Linguistics.
- Harald Baayen and Rochelle Lieber. 1996. Word frequency distributions and lexical semantics. *Computers and the Humanities*, 30(4):281–291.
- Harald Baayen, 1992. *Quantitative Aspects of Morphological Productivity*, pages 109–149. Springer Netherlands, Dordrecht.
- Harald Baayen, 1993. *On Frequency, Transparency and Productivity*, pages 181–208. Springer Netherlands, Dordrecht.
- Francis Bond and Ryan Foster. 2013. Linking and extending an open multilingual WordNet. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1352–1362.
- Francis Bond and Kyonghee Paik. 2012. A survey of WordNets and their licenses. In *Proceedings of the 6th Global WordNet Conference (GWC 2012)*, pages 64–71.
- Peter F. Brown, John Cocke, Stephen A. Della Pietra, Vincent J. Della Pietra, Fredrick Jelinek, John D. Lafferty, Robert L. Mercer, and Paul S. Roossin. 1990. A statistical approach to machine translation. *Computational Linguistics*, 16(2):79–85.
- Joan L. Bybee. 1985. *Morphology: A study of the relation between meaning and form*, volume 9. John Benjamins Publishing.
- Joan Bybee. 1995a. Diachronic and typological properties of morphology and their implications for representation. *Morphological Aspects of Language Processing*, pages 225–246.
- Joan Bybee. 1995b. Regular morphology and the lexicon. *Language and Cognitive Processes*, 10(5):425–455.

- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *Proceedings of ICLR 2018*.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, David Yarowsky, Jason Eisner, and Mans Hulden. 2016. The SIGMORPHON 2016 shared Task—Morphological reinflection. In *Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 10–22. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqi, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 2017. CoNLL-SIGMORPHON 2017 shared task: Universal morphological reinflection in 52 languages. In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 1–30.
- Paula Czarnowska, Sebastian Ruder, Edouard Grave, Ryan Cotterell, and Ann Copestake. 2019. Don’t forget the long tail! A comprehensive analysis of morphological generalization in bilingual lexicon induction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
- Georgiana Dinu, Angeliki Lazaridou, and Marco Baroni. 2015. Improving zero-shot learning by mitigating the hubness problem. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015), Workshop Track*.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Pascale Fung and Yuen Yee Lo. 1998. Translating unknown words using nonparallel, comparable texts. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics (ACL-COLING98)*.
- Pascale Fung. 1998. A statistical view on bilingual lexicon extraction: From parallel corpora to non-parallel corpora. In David Farwell, Laurie Gerber, and Eduard Hovy, editors, *Machine Translation and the Information Soup*, pages 1–17, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Goran Glavaš, Robert Litschko, Sebastian Ruder, and Ivan Vulić. 2019. How to (properly) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions. In *Proceedings of ACL 2019*.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Martin Haspelmath and Andrea Sims. 2013. *Understanding morphology*. Routledge.
- Jennifer Hay. 2001. Lexical frequency in morphology: Is everything relative? *Linguistics*, 39(6; ISSU 376):1041–1070.
- Ray Jackendoff. 1975. Morphological and semantic regularities in the lexicon. *Language*, pages 639–671.
- Yova Kementchedjhiava, Mareike Hartmann, and Anders Søgaard. 2019. Lost in evaluation: Misleading benchmarks for bilingual dictionary induction. 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 3336–3341. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Christo Kirov, John Sylak-Glassman, Roger Que, and David Yarowsky. 2016. Very-large scale parsing and normalization of Wiktionary morphological paradigms. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, Helene Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France, may. European Language Resources Association (ELRA).
- Philipp Koehn and Kevin Knight. 2002. Learning a translation lexicon from monolingual corpora. In *Proceedings of the ACL-02 workshop on Unsupervised lexical acquisition*, pages 9–16.
- Julian Kupiec. 1993. An algorithm for finding noun phrase correspondences in bilingual corpora. In *Proceedings of the 31st Annual Meeting on Association for Computational Linguistics*, ACL 93, page 1722, USA. Association for Computational Linguistics.

- Angeliki Lazaridou, Georgiana Dinu, and Marco Baroni. 2015. Hubness and pollution: Delving into cross-space mapping for zero-shot learning. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 270–280.
- Arya D. McCarthy, Ekaterina Vylomova, Shijie Wu, Chaitanya Malaviya, Lawrence Wolf-Sonkin, Garrett Nicolai, Christo Kirov, Miikka Silfverberg, Sebastian J. Mielke, Jeffrey Heinz, Ryan Cotterell, and Mans Hulden. 2019. The SIGMORPHON 2019 shared task: Morphological analysis in context and cross-lingual transfer for inflection. In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 229–244. Association for Computational Linguistics.
- Arya D. McCarthy, Christo Kirov, Matteo Grella, Amrit Nidhi, Patrick Xia, Kyle Gorman, Ekaterina Vylomova, Sabrina J. Mielke, Garrett Nicolai, Miikka Silfverberg, Timofey Arkhangelskiy, Nataly Krizhanovsky, Andrew Krizhanovsky, Elena Klyachko, Alexey Sorokin, John Mansfield, Valts Ernštreits, Yuval Pinter, Cassandra L. Jacobs, Ryan Cotterell, Mans Hulden, and David Yarowsky. 2020. UniMorph 3.0: Universal Morphology. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3922–3931, Marseille, France, May. European Language Resources Association.
- James L. McClelland, David E. Rumelhart, PDP Research Group, et al. 1987. *Parallel distributed processing*, volume 2. MIT Press.
- Antonio Valerio Miceli Barone. 2016. Towards cross-lingual distributed representations without parallel text trained with adversarial autoencoders. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 121–126, Berlin, Germany, August. Association for Computational Linguistics.
- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation. *CoRR*, abs/1309.4168.
- Ndapa Nakashole and Raphael Flauger. 2018. Characterizing departures from linearity in word translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 221–227. Association for Computational Linguistics.
- Aitor Ormazabal, Mikel Artetxe, Gorika Labaka, Aitor Soroa, and Eneko Agirre. 2019. Analyzing the limitations of cross-lingual word embedding mappings. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Barun Patra, Joel Ruben Antony Moniz, Sarthak Garg, Matthew R. Gormley, and Graham Neubig. 2019. Bilingual lexicon induction with semi-supervision in non-isometric embedding spaces. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 184–193. Association for Computational Linguistics.
- Steven Pinker and Alan Prince. 1994. Regular and irregular morphology and the psychological status of rules of grammar. *The Reality of Linguistic Rules*, 321:51.
- Steven Pinker. 1998. Words and rules. *Lingua*, 106(1-4):219–242.
- Sandeep Prasada and Steven Pinker. 1993. Generalisation of regular and irregular morphological patterns. *Language and Cognitive Processes*, 8(1):1–56.
- Reinhard Rapp. 1995. Identifying word translations in non-parallel texts. *ACL '95: Proceedings of the 33rd annual meeting on Association for Computational Linguistics*, pages 320–322.
- Parker Riley and Daniel Gildea. 2018. Orthographic features for bilingual lexicon induction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 390–394.
- Sebastian Ruder, Ryan Cotterell, Yova Kementchedjheva, and Anders Søgaard. 2018. A discriminative latent-variable model for bilingual lexicon induction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 458–468. Association for Computational Linguistics.

- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2019. A survey of cross-lingual word embedding models. *Journal of Artificial Intelligence Research*, 65:569–631.
- Frank Smadja and Kathleen McKeown. 1994. Translating collocations for use in bilingual lexicons. In *Proceedings of the workshop on Human Language Technology*, pages 152–156. Association for Computational Linguistics.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. 2018. On the limitations of unsupervised bilingual dictionary induction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788. Association for Computational Linguistics.
- Stephan Vogel, Hermann Ney, and Christoph Tillmann. 1996. HMM-based word alignment in statistical translation. In *COLING 1996 Volume 2: The 16th International Conference on Computational Linguistics*.
- Ivan Vulić, Goran Glavaš, Roi Reichart, and Anna Korhonen. 2019. Do we really need fully unsupervised cross-lingual embeddings? *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
- Shijie Wu, Ryan Cotterell, and Timothy J. O’Donnell. 2019a. Morphological irregularity correlates with frequency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5117–5126.
- Shijie Wu, Pamela Shapiro, and Ryan Cotterell. 2019b. Exact hard monotonic attention for character-level transduction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5117–5126.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011.
- Pengcheng Yang, Fuli Luo, Peng Chen, Tianyu Liu, and Xu Sun. 2019. MAAM: A morphology-aware alignment model for unsupervised bilingual lexicon induction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3190–3196. Association for Computational Linguistics.
- Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. Adversarial training for unsupervised bilingual lexicon induction. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1959–1970.

Appendix A Hyperparameter settings

Model	encoder layers	decoder layers	encoder dim.	decoder dim.	embedding dim.	init. learning rate	batch size	dropout
Inflector	2	1	400	400	200	0.001	20	0.4
Analyzer	2	1	100	100	50	0.001	20	0.4

	fra-ita	fra-spa	ita-fra	ita-spa	spa-fra	spa-ita	eng-ita	eng-spa	spa-eng	ita-eng
batch size	24	24	24	24	24	24	24	24	24	24
α	15	5	15	10	15	10	10	10	10	10
init. learn. rate	0.05	0.025	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05

Table 3: The optimization and hyperparameter settings for the analysis and inflection modules across all language pairs (above) and for the translator module (below). α is the orthogonal regularization weight. During training we divide the values of α by the batch size to match the weight of the training loss (the losses are averaged across observations for each mini batch).

Appendix B Transducer Performance

Task	FRA-ITA	FRA-SPA	ITA-FRA	ITA-SPA	SPA-FRA	SPA-ITA
Inflection	99.1	97.4	94.2	97.9	93.6	99.1
Analysis	89.5	89.7	91.1	91.0	91.1	91.5

Table 4: Transducers’ accuracy on the development split of the dictionaries of Czarnowska et al. (2019), when trained on UniMorph entries (excluding those present in the test or development splits of the dictionaries).

Table 4 displays the accuracy of the analysis and inflection components on the development set of the dictionaries of Czarnowska et al. (2019). For most language pairs the performance of the inflectors is in the high 90s, while the performance of the analyzers averages at 90.7%. Indeed, the word-type analysis is inherently more difficult than inflection, as it is less deterministic: in most cases there exist more than one correct lemma-tag output. Note that this task is different from the word-token analysis, where the *context* of the analyzed inflection is known. In addition, in contrast to inflection where the model is given the morpho-syntactic tag, the analyzer has no information about the type of word it is handling. In particular, in an additional line of experimentation, we found that adding a POS to the analyzed source form notably improves the performance. Similarly, training and evaluating on forms associated with a single POS also leads to better accuracy.