

A Call for More Rigor in Unsupervised Cross-lingual Learning

Mikel Artetxe^{†*}, Sebastian Ruder^{†*}, Dani Yogatama[‡], Gorka Labaka[†], Eneko Agirre[†]

[†]HiTZ Center, University of the Basque Country (UPV/EHU)

[‡]DeepMind

{mikel.artetxe, gorka.labaka, e.agirre}@ehu.eus

{ruder, dyogatama}@google.com

Abstract

We review motivations, definition, approaches, and methodology for unsupervised cross-lingual learning and call for a more rigorous position in each of them. An existing rationale for such research is based on the lack of parallel data for many of the world’s languages. However, we argue that a scenario without *any* parallel data and abundant monolingual data is unrealistic in practice. We also discuss different training signals that have been used in previous work, which depart from the pure unsupervised setting. We then describe common methodological issues in tuning and evaluation of unsupervised cross-lingual models and present best practices. Finally, we provide a unified outlook for different types of research in this area (i.e., cross-lingual word embeddings, deep multilingual pretraining, and unsupervised machine translation) and argue for comparable evaluation of these models.

1 Introduction

The study of the connection among human languages has contributed to major discoveries including the evolution of languages, the reconstruction of proto-languages, and an understanding of language universals (Eco and Fentress, 1995). In natural language processing, the main promise of multilingual learning is to bridge the digital language divide, to enable access to information and technology for the world’s 6,900 languages (Ruder et al., 2019). For the purpose of this paper, we define “*multilingual learning*” as learning a common model for two or more languages from raw text, without any downstream task labels. Common use cases include translation as well as pretraining multilingual representations. We will use the term interchangeably with “*cross-lingual learning*”.

Recent work in this direction has increasingly focused on purely unsupervised cross-lingual learning (UCL)—i.e., cross-lingual learning without any parallel signal across the languages. We provide an overview in §2. Such work has been motivated by the apparent dearth of parallel data for most of the world’s languages. In particular, previous work has noted that “*data encoding cross-lingual equivalence is often expensive to obtain*” (Zhang et al., 2017a) whereas “*monolingual data is much easier to find*” (Lample et al., 2018a). Overall, it has been argued that unsupervised cross-lingual learning “*opens up opportunities for the processing of extremely low-resource languages and domains that lack parallel data completely*” (Zhang et al., 2017a).

We challenge this narrative and argue that the scenario of no parallel data *and* sufficient monolingual data is unrealistic and not reflected in the real world (§3.1). Nevertheless, UCL is an important research direction and we advocate for its study based on an inherent scientific interest (to better understand and make progress on general language understanding), usefulness as a lab setting, and simplicity (§3.2).

Unsupervised cross-lingual learning permits no supervisory signal by definition. However, previous work implicitly includes monolingual and cross-lingual signals that constitute a departure from the pure setting. We review existing training signals as well as other signals that may be of interest for future study (§4). We then discuss methodological issues in UCL (e.g., validation, hyperparameter tuning) and propose best evaluation practices (§5). Finally, we provide a unified outlook of established research areas (cross-lingual word embeddings, deep multilingual models and unsupervised machine translation) in UCL (§6), and conclude with a summary of our recommendations (§7).

*Equal contribution.

2 Background

In this section, we briefly review existing work on UCL, covering cross-lingual word embeddings (§2.1), deep multilingual pre-training (§2.2), and unsupervised machine translation (§2.3).

2.1 Cross-lingual word embeddings

Cross-lingual word embedding methods traditionally relied on parallel corpora (Gouws et al., 2015; Luong et al., 2015). Nonetheless, the amount of supervision required was greatly reduced via cross-lingual word embedding mappings, which work by separately learning monolingual word embeddings in each language and mapping them into a shared space through a linear transformation. Early work required a bilingual dictionary to learn such a transformation (Mikolov et al., 2013a; Faruqui and Dyer, 2014). This requirement was later reduced with self-learning (Artetxe et al., 2017), and ultimately removed via unsupervised initialization heuristics (Artetxe et al., 2018a; Hoshen and Wolf, 2018) and adversarial learning (Zhang et al., 2017a; Conneau et al., 2018a). Finally, several recent methods have formulated cross-lingual embedding alignment as an optimal transport problem (Zhang et al., 2017b; Grave et al., 2019; Alvarez-Melis and Jaakkola, 2018).

2.2 Deep multilingual pretraining

Following the success in learning shallow word embeddings (Mikolov et al., 2013b; Pennington et al., 2014), there has been an increasing interest in learning contextual word representations (Dai and Le, 2015; Peters et al., 2018; Howard and Ruder, 2018). Recent research has been dominated by BERT (Devlin et al., 2019), which uses a bidirectional transformer encoder trained on masked language modeling and next sentence prediction, which led to impressive gains on various downstream tasks.

While the above approaches are limited to a single language, a multilingual extension of BERT (mBERT) has been shown to also be effective at learning cross-lingual representations in an unsupervised way.¹ The main idea is to combine monolingual corpora in different languages, upsampling those with less data, and training a regular BERT model on the combined data. Conneau and Lample (2019) follow a similar approach but perform a more thorough evaluation and report substantially

¹<https://github.com/google-research/bert/blob/master/multilingual.md>

stronger results,² which was further scaled up by Conneau et al. (2019). Several recent studies (Wu and Dredze, 2019; Pires et al., 2019; Artetxe et al., 2020b; Wu et al., 2019) analyze mBERT to get a better understanding of its capabilities.

2.3 Unsupervised machine translation

Early attempts to build machine translation systems using monolingual data alone go back to statistical decipherment (Ravi and Knight, 2011; Dou and Knight, 2012, 2013). However, this approach was only shown to work in limited settings, and the first convincing results on standard benchmarks were achieved by Artetxe et al. (2018c) and Lample et al. (2018a) on unsupervised Neural Machine Translation (NMT). Both approaches rely on cross-lingual word embeddings to initialize a shared encoder, and train it in conjunction with the decoder using a combination of denoising autoencoding, back-translation, and optionally adversarial learning.

Subsequent work adapted these principles to unsupervised phrase-based Statistical Machine Translation (SMT), obtaining large improvements over the original NMT-based systems (Lample et al., 2018b; Artetxe et al., 2018b). This alternative approach uses cross-lingual n -gram embeddings to build an initial phrase table, which is combined with an n -gram language model and a distortion model, and further refined through iterative back-translation. There have been several follow-up attempts to combine NMT and SMT based approaches (Marie and Fujita, 2018; Ren et al., 2019; Artetxe et al., 2019b). More recently, Conneau and Lample (2019), Song et al. (2019) and Liu et al. (2020) obtain strong results using deep multilingual pretraining rather than cross-lingual word embeddings to initialize unsupervised NMT systems.

3 Motivating fully unsupervised learning

In this section, we challenge the narrative of motivating UCL based on a lack of parallel resources. We argue that the strict unsupervised scenario cannot be motivated from an immediate practical perspective, and elucidate what we believe should be the true goals of this research direction.

²The full version of their model (XLM) requires parallel corpora for their translation language modeling objective, but the authors also explore an unsupervised variant using masked language modeling alone.

3.1 How practical is the strict unsupervised scenario?

Monolingual resources subsume parallel resources. For instance, each side of a parallel corpus effectively serves as a monolingual corpus. From this argument, it follows that monolingual data is cheaper to obtain than parallel data, so unsupervised cross-lingual learning should in principle be more generally applicable than supervised learning.

However, we argue that the common claim that the requirement for **parallel data** “*may not be met for many language pairs in the real world*” (Xu et al., 2018) is largely inaccurate. For instance, the JW300 parallel corpus covers 343 languages with around 100,000 parallel sentences per language pair on average (Agić and Vulić, 2019), and the multilingual Bible corpus collected by Mayer and Cysouw (2014) covers 837 language varieties (each with a unique ISO 639-3 code). Moreover, the PanLex project aims to collect multilingual lexica for all human languages in the world, and already covers 6,854 language varieties with at least 20 lexemes, 2,364 with at least 200 lexemes, and 369 with at least 2,000 lexemes (Kamholz et al., 2014). While 20 or 200 lexemes might seem insufficient, weakly supervised cross-lingual word embedding methods already proved effective with as little as 25 word pairs (Artetxe et al., 2017). More recent methods have focused on completely removing this weak supervision (Conneau et al., 2018a; Artetxe et al., 2018a), which can hardly be justified from a practical perspective given the existence of such resources and additional training signals stemming from a (partially) shared script (§4.2). Finally, given the availability of sufficient monolingual data, noisy parallel data can often be obtained by mining bitext (Schwenk et al., 2019a,b).

In addition, large **monolingual data** is difficult to obtain for low-resource languages. For instance, recent work on cross-lingual word embeddings has mostly used Wikipedia as its source for monolingual corpora (Gouws et al., 2015; Vulić and Korhonen, 2016; Conneau et al., 2018a). However, as of November 2019, Wikipedia exists in *only* 307 languages³ of which nearly half have less than 10,000 articles. While one could hope to overcome this by taking the entire web as a corpus, as facilitated by Common Crawl⁴ and similar initiatives, this is not

³https://en.wikipedia.org/wiki/List_of_Wikipedias

⁴<https://commoncrawl.org/>

always feasible for low-resource languages. First, the presence of less resourced languages on the web is very limited, with only a few hundred languages recognized as being used in websites.⁵ This situation is further complicated by the limited coverage of existing tools such as language detectors (Buck et al., 2014; Grave et al., 2018), which only cover a few hundred languages. Alternatively, speech could also serve as a source of monolingual data (e.g., by recording public radio stations). However, this is an unexplored direction within UCL, and collecting, processing and effectively capitalizing on speech data is far from trivial, particularly for low-resource languages.

All in all, we conclude that the alleged scenario involving no parallel data *and* sufficient monolingual data is not met in the real world in the terms explored by recent UCL research. Needless to say, effectively exploiting unlabeled data is important in any low-resource setting. However, refusing to use an informative training signal—which parallel data is—when it does indeed exist, cannot be justified from a practical perspective if one’s goal is to build the strongest possible model. For this reason, we believe that semi-supervised learning is a more suitable paradigm for truly low-resource languages, and UCL should not be motivated from an immediate practical perspective.

3.2 A scientific motivation

Despite not being an entirely realistic setup, we believe that UCL is an important research direction for the reasons we discuss below.

Inherent scientific interest. The extent to which two languages can be aligned based on independent samples—without any cross-lingual signal—is an open and scientifically relevant problem *per se*. In fact, it is not entirely obvious that UCL should be possible at all, as humans would certainly struggle to align two unknown languages without any grounding. Exploring the limits of UCL could help to understand the limits of the principles that the corresponding methods are based on, such as the distributional hypothesis. Moreover, this research line could bring new insights into the properties and inner workings of both language acquisition and the underlying computational models that ultimately make UCL possible. Finally, such methods may be useful in areas where supervision is impos-

⁵https://w3techs.com/technologies/overview/content_language

sible to obtain, such as when dealing with unknown or even non-human languages.

Useful as a lab setting. The strict unsupervised scenario, although not practical, allows us to isolate and better study the use of monolingual corpora for cross-lingual learning. We believe lessons learned in this setting can be useful in the more practical semi-supervised scenario. In a similar vein, monolingual language models, although hardly useful on their own, have contributed to large improvements in other tasks. From a research methodology perspective, unsupervised systems also set a competitive baseline, which any semi-supervised method should improve upon.

Simplicity as a value. As we discussed previously, refusing to use an informative training signal when it does exist can hardly be beneficial, so we should not expect UCL to perform better than semi-supervised learning. However, simplicity is a value in its own right. Unsupervised approaches could be preferable to their semi-supervised counterparts if the performance gap between them is small enough. For instance, unsupervised cross-lingual embedding methods have been reported to be competitive with their semi-supervised counterparts in certain settings (Glavaš et al., 2019), while being easier to use in the sense that they do not require a bilingual dictionary.

4 What does *unsupervised* mean?

In its most general sense, unsupervised cross-lingual learning can be seen as referring to any method relying *exclusively* on monolingual text data in two or more languages. However, there are different training signals—stemming from common assumptions and varying amounts of linguistic knowledge—that one can potentially exploit under such a regime. This has led to an inconsistent use of this term in the literature. In this section, we categorize different training signals available both from a monolingual and a cross-lingual perspective and discuss additional scenarios enabled by multiple languages.

4.1 Monolingual training signals

From a computational perspective, text is modeled as a sequence of discrete symbols. In UCL, the training data consists of a set of such sequences in each of the languages. In principle, without any knowledge about the languages, one would have no

prior information of the nature of such sequences or the possible relations between them. In practice, however, sets of sequences are assumed to be independent, and existing work differs whether they assume document-level sequences (Conneau and Lample, 2019) or sentence-level sequences (Artetxe et al., 2018c; Lample et al., 2018a).

Nature of atomic symbols. A more important consideration is the nature of the atomic symbols in such sequences. To the best of our knowledge, previous work assumes some form of word segmentation or tokenization (e.g., splitting by whitespaces or punctuation marks). Early work on cross-lingual word embeddings considered such tokens as atomic units. However, more recent work (Hoshen and Wolf, 2018; Glavaš et al., 2019) has primarily used fastText embeddings (Bojanowski et al., 2017) which incorporate subword information into the embedding learning, although the vocabulary is still defined at the token level. In addition, there have also been approaches that incorporate character-level information into the alignment learning itself (Heyman et al., 2017; Riley and Gildea, 2018). In contrast, most work on contextual word embeddings and unsupervised machine translation operates with a subword vocabulary (Devlin et al., 2019; Conneau and Lample, 2019).

While the above distinction might seem irrelevant from a practical perspective, we think that it is important from a more fundamental point of view (e.g. in relation to the distributional hypothesis as discussed in §3.2). Moreover, some of the underlying assumptions might not generalize to different writing systems (e.g. logographic instead of alphabetic). For instance, subword tokenization has been shown to perform poorly on reduplicated words (Vania and Lopez, 2017). In relation to that, one could also consider the text in each language as a stream of discrete character-like symbols without any notion of tokenization. Such a *tabula rasa* approach is potentially applicable to any arbitrary language, even when its writing system is not known, but has so far only been explored for a limited number of languages in a monolingual setting (Hahn and Baroni, 2019).

Linguistic information. Finally, one can exploit additional linguistic knowledge through linguistic analysis such as lemmatization, part-of-speech tagging, or syntactic parsing. For instance, before the advent of unsupervised NMT, statistical deci-

phering was already shown to benefit from incorporating syntactic dependency relations (Dou and Knight, 2013). For other tasks such as unsupervised POS tagging (Snyder et al., 2008), monolingual tag dictionaries have been used. While such approaches could still be considered unsupervised from a cross-lingual perspective, we argue that the interest of this research direction is greatly limited by two factors: (i) from a theoretical perspective, it assumes some fundamental knowledge that is not directly inferred from the raw monolingual corpora; and (ii) from a more practical perspective, it is not reasonable to assume that such resources are available in the less resourced settings where this research direction has more potential for impact.

4.2 Cross-lingual training signals

Pure UCL should not use any cross-lingual signal by definition. When we view text as a sequence of discrete atomic symbols (either characters or tokens), a strict interpretation of this principle would consider the set of atomic symbols in different languages to be disjoint, without prior knowledge of the relationship between them.

Needless to say, any form of learning requires making assumptions, as one needs some criterion to prefer one mapping over another. In the case of UCL, such assumptions stem from the structural similarity across languages (e.g. semantically equivalent words in different languages are assumed to occur in similar contexts). In practice, these assumptions weaken as the distribution of the datasets diverges, and some UCL models have been reported to break under a domain shift (Søgaard et al., 2018; Guzmán et al., 2019; Marchisio et al., 2020). Similarly, approaches that leverage linguistic features such as syntactic dependencies may assume that these are similar across languages.

In addition, one can also assume that the sets of symbols that are used to represent different languages have some commonalities. This departs from the strict definition of UCL above, establishing some prior connections between the sets of symbols in different languages. Such an assumption is reasonable from a practical perspective, as there are a few scripts (e.g. Latin, Arabic or Cyrillic) that cover a large fraction of languages. Moreover, even when two languages use different writing systems or scripts, there are often certain elements that are still shared (e.g. Arabic numerals, named entities written in a foreign script, URLs, certain punctua-

tion marks, etc.). In relation to that, several models have relied on identically spelled words (Artetxe et al., 2017; Smith et al., 2017; Søgaard et al., 2018) or string-level similarity across languages (Riley and Gildea, 2018; Artetxe et al., 2019b) as training signals. Other methods use a joint subword vocabulary for all languages, indirectly exploiting the commonalities in their writing system (Lample et al., 2018b; Conneau and Lample, 2019).

However, past work greatly differs on the nature and relevance that is attributed to such a training signal. The reliance on identically spelled words has been considered as a weak form of supervision in the cross-lingual word embedding literature (Søgaard et al., 2018; Ruder et al., 2018), and significant effort has been put into developing strictly unsupervised methods that do not rely on such signal (Conneau et al., 2018a). In contrast, the unsupervised machine translation literature has not paid much attention to this factor, and has often relied on identical words (Artetxe et al., 2018c), string-level similarity (Artetxe et al., 2019b), or a joint subword vocabulary (Lample et al., 2018b; Conneau and Lample, 2019) under the unsupervised umbrella. The same is true for unsupervised deep multilingual pretraining, where a shared subword vocabulary has been a common component (Pires et al., 2019; Conneau and Lample, 2019), although recent work shows that it is not important to share vocabulary across languages (Artetxe et al., 2020b; Wu et al., 2019).

Our position is that making assumptions on linguistics universals is acceptable and ultimately necessary for UCL. However, we believe that any connection stemming from a (partly) shared writing system belongs to a different category, and should be considered a separate cross-lingual signal. Our rationale is that a given writing system pertains to a specific form to encode a language, but cannot be considered to be part of the language itself.⁶

4.3 Multilinguality

While most work in unsupervised cross-lingual learning considers two languages at a time, there have recently been some attempts to extend these methods to multiple languages (Duong et al., 2017; Chen and Cardie, 2018; Heyman et al., 2019), and most work on unsupervised cross-lingual pretraining is multilingual (Pires et al., 2019; Conneau

⁶As a matter of fact, languages existed well before writing was invented, and a given language can have different writing systems or new ones can be designed.

Monolingual signal	Cross-lingual signal
Sequence of symbols	Shared writing system
Sets of sentences/documents	Identical words
Tokens/subwords	String similarity
Linguistic analysis	

Table 1: Different types of monolingual and cross-lingual signals that have been used for unsupervised cross-lingual learning, ordered roughly from least to most linguistic knowledge (top to bottom).

and Lample, 2019). When considering parallel data across a subset of the language pairs, multilinguality gives rise to additional scenarios. For instance, the scenario where two languages have no parallel data between each other but are well connected through a third (pivot) language has been explored by several authors in the context of machine translation (Cheng et al., 2016; Chen et al., 2017). However, given that the languages in question are still indirectly connected through parallel data, this scenario does not fall within the *unsupervised* category, and is instead commonly known as *zero-resource* machine translation.

An alternative scenario explored in the contemporary work of Liu et al. (2020) is where a set of languages are connected through parallel data, and there is a separate language with monolingual data only. We argue that, when it comes to the isolated language, such a scenario should still be considered as UCL, as it does not rely on any parallel data for that particular language nor does it assume any previous knowledge of it. This scenario is easy to justify from a practical perspective given the abundance of parallel data for high-resource languages, and can also be interesting from a more theoretical point of view. This way, rather than considering two unknown languages, this alternative scenario would assume some knowledge of how one particular language is connected to other languages, and attempt to align it to a separate unknown language.

4.4 Discussion

As discussed throughout the section, there are different training signals that we can exploit depending on the available resources of the languages involved and the assumptions made regarding their writing system, which are summarized in Table 1. Many of these signals are not specific to work on UCL but have been observed in the past in allegedly language-independent NLP approaches, as discussed by Bender (2011). Others, such as a re-

liance on subwords or shared symbols are more recent phenomena.

While we do not aim to open a terminological debate on what UCL encompasses, we advocate for future work being more aware and explicit about the monolingual and cross-lingual signals they employ, what assumptions they make (e.g. regarding the writing system), and the extent to which these generalize to other languages.

In particular, we argue that it is critical to consider the assumptions made by different methods when comparing their results. Otherwise the blind chase for state-of-the-art performance may benefit models making stronger assumptions and exploiting all available training signals, which could ultimately conflict with the eminently scientific motivation of this research area (see §3.2).

5 Methodological issues

In this section, we describe methodological issues that are commonly encountered when training and evaluating unsupervised cross-lingual models and propose measures to ameliorate them.

5.1 Validation and hyperparameter tuning

In conventional supervised or semi-supervised settings, we use a separate validation set for development and hyperparameter tuning. However, this becomes tricky in unsupervised cross-lingual learning, where we ideally should not use any parallel data other than for testing purposes.

Previous work has not paid much attention to this aspect, and different methods are evaluated with different validation schemes. For instance, Artetxe et al. (2018b,c) use a **separate language pair** with a parallel validation set to make all development and hyperparameter decisions. They test their final system on other language pairs without any parallel data. This approach has the advantage of being strictly unsupervised with respect to the test language pairs, but the optimal hyperparameter choice might not necessarily transfer well across languages. In contrast, Conneau et al. (2018a) and Lample et al. (2018a) propose an **unsupervised validation criterion** that is defined over monolingual data and shown to correlate well with test performance. This enables systematic tuning on the language pair of interest, but still requires parallel data to guide the development of the unsupervised validation criterion itself. A **parallel validation set** has also been used for systematic tuning in

the context of unsupervised machine translation (Marie and Fujita, 2018; Marie et al., 2019; Stojanovski et al., 2019). While this is motivated as a way to abstract away the issue of unsupervised tuning—which the authors consider to be an open problem—we argue that any systematic use of parallel data should not be considered UCL. Finally, previous work often does not report the validation scheme used. In particular, unsupervised cross-lingual word embedding methods have almost exclusively been evaluated on bilingual lexicons that do not have a validation set, and presumably use the **test set** to guide development to some extent.

Our position is that a completely blind development model without any parallel data is unrealistic. Some cross-lingual signals to guide development are always needed. However, this factor should be carefully controlled and reported with the necessary rigor as a part of the experimental design. We advocate for using one language pair for development and evaluating on others when possible. If parallel data in the target language pair is used, the test set should be kept blind to avoid overfitting, and a separate validation should be used. In any case, we argue that the use of parallel data in the target language pair should be minimized if not completely avoided, and it should under no circumstances be used for extensive tuning. Instead, we recommend to use unsupervised validation criteria for systematic tuning in the target language.

5.2 Evaluation practices

We argue that there are also several issues with common evaluation practices in UCL.

Evaluation on favorable conditions. Most work on UCL has focused on relatively close languages with large amounts of high-quality parallel corpora from similar domains. Only recently have approaches considered more diverse languages as well as language pairs that do not involve English (Glavaš et al., 2019; Vulić et al., 2019), and some existing methods have been shown to completely break in less favorable conditions (Guzmán et al., 2019; Marchisio et al., 2020). In addition, most approaches have focused on learning from similar domains, often involving Wikipedia and news corpora, which are unlikely to be available for low-resource languages. We believe that future work should pay more attention to the effect of the typology and linguistic distance of the languages involved, as well as the size, noise and domain

similarity of the training data used.

Over-reliance on translation tasks. Most work on UCL focuses on translation tasks, either at the word level (where the problem is known as *bilingual lexicon induction*) or at the sentence level (where the problem is known as *unsupervised machine translation*). While translation can be seen as the ultimate application of cross-lingual learning and has a strong practical interest on its own, it only evaluates a particular facet of a model’s cross-lingual generalization ability. In relation to that, Glavaš et al. (2019) showed that bilingual lexicon induction performance does not always correlate well with downstream tasks. In particular, they observe that some mapping methods that are specifically designed for bilingual lexicon induction perform poorly on other tasks, showing the risk of relying excessively on translation benchmarks for evaluating cross-lingual models.

Moreover, existing translation benchmarks have been shown to have several issues on their own. In particular, bilingual lexicon induction datasets have been reported to misrepresent morphological variations, overly focus on named entities and frequent words, and have pervasive gaps in the gold-standard targets (Czarnowska et al., 2019; Kementchedjieva et al., 2019). More generally, most of these datasets are limited to relatively close languages and comparable corpora.

Lack of an established cross-lingual benchmark. At the same time, there is no *de facto* standard benchmark to evaluate cross-lingual models beyond translation. Existing approaches have been evaluated in a wide variety of tasks including dependency parsing (Schuster et al., 2019), named entity recognition (Rahimi et al., 2019), sentiment analysis (Barnes et al., 2018), natural language inference (Conneau et al., 2018b), and document classification (Schwenk and Li, 2018). XNLI (Conneau et al., 2018b) and MLDoc (Schwenk and Li, 2018) are common choices, but they have their own problems: MultiNLI, the dataset from which XNLI was derived, has been shown to contain superficial cues that can be exploited (Gururangan et al., 2018), while MLDoc can be solved by keyword matching (Artetxe et al., 2020b). There are non-English counterparts for more challenging tasks such as question answering (Cui et al., 2019; Hsu et al., 2019), but these only exist for a handful of languages. More recent datasets such as XQuAD

Methodological issues	Examples
Validation and hyperparameter tuning	Systematic tuning with parallel data or on test data
Evaluation on favorable conditions	Typologically similar languages; always including English; training on the same domain
Over-reliance on translation tasks	Overfitting to bilingual lexicon induction; known issues with existing datasets
Lack of an established benchmark	Evaluation on many different tasks; problems with common tasks (MLDoc and XNLI)

Table 2: Methodological issues pertaining to validation and hyperparameter tuning and evaluation practices in current work on unsupervised cross-lingual learning.

(Artetxe et al., 2020b), MLQA (Lewis et al., 2019) and TyDi QA (Clark et al., 2020) cover a wider set of languages, but a comprehensive benchmark that evaluates multilingual representations on a diverse set of tasks—in the style of GLUE (Wang et al., 2018)—and languages has been missing until very recently. The contemporaneous XTREME (Hu et al., 2020) and XGLUE (Liang et al., 2020) benchmarks try to close this gap, but they are still restricted to languages where existing labelled data is available. Finally, an additional issue is that a large part of these benchmarks were created through translation, which was recently shown to introduce artifacts (Artetxe et al., 2020a).

We present a summary of the methodological issues discussed in Table 2.

6 Bridging the gap between unsupervised cross-lingual learning flavors

The three categories of UCL (§2) have so far been treated as separate research topics by the community. In particular, cross-lingual word embeddings have a long history (Ruder et al., 2019), while deep multilingual pretraining has emerged as a separate line of research with its own best practices and evaluation standards. At the same time, unsupervised machine translation has been considered a separate problem in its own right, where cross-lingual word embeddings and deep multilingual pretraining have just served as initialization techniques.

While each of these families have their own defining features, we believe that they share a strong connection that should be considered from a more holistic perspective. In particular, both cross-lingual word embeddings and deep mul-

tilingual pretraining share the goal of learning (sub)word representations, and essentially differ on whether such representations are static or context-dependent. Similarly, in addition to being a downstream application of the former, unsupervised machine translation can also be useful to develop other multilingual applications or learn better cross-lingual representations. This has previously been shown for *supervised* machine translation (McCann et al., 2017; Siddhant et al., 2019) and recently for bilingual lexicon induction (Artetxe et al., 2019a). In light of these connections, we call for a more holistic view of UCL, both from an experimental and theoretical perspective.

Evaluation. Most work on cross-lingual word embeddings focuses on bilingual lexicon induction. In contrast, deep multilingual pretraining has not been tested on this task, and is instead typically evaluated on zero-shot cross-lingual transfer. We think it is important to evaluate both families—cross-lingual word embeddings and deep multilingual representations—in the same conditions to better understand their strengths and weaknesses. In that regard, Artetxe et al. (2020b) recently showed that deep pretrained models are much stronger in some downstream tasks, while cross-lingual word embeddings are more efficient and sufficient for simpler tasks. However, this could partly be attributed to a particular integration strategy, and we advocate for using a common evaluation framework in future work to allow a direct comparison between the different families.

Theory. From a more theoretical perspective, it is still not well understood in what ways cross-lingual word embeddings and deep multilingual pretraining differ. While one could expect the latter to be learning higher-level multilingual abstractions, recent work suggests that deep multilingual models might mostly be learning a lexical-level alignment (Artetxe et al., 2020b). For that reason, we believe that further research is needed to understand the relation between both families of models.

7 Recommendations

To summarize, we make the following practical recommendations for future cross-lingual research:

- Be rigorous when motivating UCL. Do not present it as a practical scenario unless supported by a real use case.

- Be explicit about the monolingual and cross-lingual signals used by your approach and the assumptions it makes, and take them into considerations when comparing different models.
- Report the validation scheme used. Minimize the use of parallel data by preferring an unsupervised validation criterion and/or using only one language for development. Always keep the test set blind.
- Pay attention to the conditions in which you evaluate your model. Consider the impact of typology, linguistic distance, and the domain similarity, size and noise of the training data. Be aware of known issues with common benchmarks, and favor evaluation on a diverse set of tasks.
- Keep a holistic view of UCL, including cross-lingual word embeddings, deep multilingual pretraining and unsupervised machine translation. To the extent possible, favor a common evaluation framework for these different families.

8 Conclusions

In this position paper, we review the status quo of unsupervised cross-lingual learning—a relatively recent field. UCL is typically motivated by the lack of cross-lingual signal for many of the world’s languages, but available resources indicate that a scenario with no parallel data and sufficient monolingual data is not realistic. Instead, we advocate for the importance of UCL for scientific reasons.

We also discuss different monolingual and cross-lingual training signals that have been used in the past, and advocate for carefully reporting them to enable a meaningful comparison across different approaches. In addition, we describe methodological issues related to the unsupervised setting and propose measures to ameliorate them. Finally, we discuss connections between cross-lingual word embeddings, deep multilingual pre-training, and unsupervised machine translation, calling for an evaluation on an equal footing.

We hope that this position paper will serve to strengthen research in UCL, providing a more rigorous look at the motivation, definition, and methodology. In light of the unprecedented growth of our field in recent times, we believe that it is essential to establish a rigorous foundation connecting past and present research, and an evaluation protocol that

carefully controls for the use of parallel data and assesses models in diverse, challenging settings.

Acknowledgments

This research was partially funded by a Facebook Fellowship, the Basque Government excellence research group (IT1343-19), the Spanish MINECO (UnsupMT TIN2017-91692-EXP MCIU/AEI/FEDER, UE) and Project BigKnowledge (Ayudas Fundación BBVA a equipos de investigación científica 2018).

References

- Željko Agić and Ivan Vulić. 2019. [JW300: A wide-coverage parallel corpus for low-resource languages](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3204–3210, Florence, Italy. Association for Computational Linguistics.
- David Alvarez-Melis and Tommi Jaakkola. 2018. [Gromov-wasserstein alignment of word embedding spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1881–1890, Brussels, Belgium. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. [Learning bilingual word embeddings with \(almost\) no bilingual data](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Vancouver, Canada. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018a. [A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798, Melbourne, Australia. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018b. [Unsupervised statistical machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3632–3642, Brussels, Belgium. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2019a. [Bilingual lexicon induction through unsupervised machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5002–5007, Florence, Italy. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2019b. [An effective approach to unsupervised machine translation](#). In *Proceedings of the 57th Annual*

- Meeting of the Association for Computational Linguistics*, pages 194–203, Florence, Italy. Association for Computational Linguistics.
- Mikel Artetxe, Gorika Labaka, and Eneko Agirre. 2020a. [Translation artifacts in cross-lingual transfer learning](#). *arXiv preprint arXiv:2004.04721*.
- Mikel Artetxe, Gorika Labaka, Eneko Agirre, and Kyunghyun Cho. 2018c. [Unsupervised neural machine translation](#). In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020b. [On the Cross-lingual Transferability of Monolingual Representations](#). In *Proceedings of ACL 2020*.
- Jeremy Barnes, Roman Klinger, and Sabine Schulte im Walde. 2018. [Bilingual sentiment embeddings: Joint projection of sentiment across languages](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2483–2493, Melbourne, Australia. Association for Computational Linguistics.
- Emily M. Bender. 2011. [On Achieving and Evaluating Language-Independence in NLP](#). *Linguistic Issues in Language Technology*, 6(3):1–26.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Christian Buck, Kenneth Heafield, and Bas van Ooyen. 2014. [N-gram counts and language models from the common crawl](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3579–3584, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Xilun Chen and Claire Cardie. 2018. [Unsupervised multilingual word embeddings](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 261–270, Brussels, Belgium. Association for Computational Linguistics.
- Yun Chen, Yang Liu, Yong Cheng, and Victor O.K. Li. 2017. [A teacher-student framework for zero-resource neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1925–1935, Vancouver, Canada. Association for Computational Linguistics.
- Yong Cheng, Yang Liu, Qian Yang, Maosong Sun, and Wei Xu. 2016. [Neural machine translation with pivot languages](#). *arXiv preprint arXiv:1611.04928*.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages](#). *Transactions of the Association for Computational Linguistics*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). *arXiv preprint arXiv:1911.02116*.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems 32*, pages 7057–7067.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018a. [Word translation without parallel data](#). In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018b. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2019. [Cross-lingual machine reading comprehension](#). 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 1586–1595, Hong Kong, China. Association for Computational Linguistics.
- Paula Czarnecka, 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 and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 973–982, Hong Kong, China. Association for Computational Linguistics.
- Andrew M. Dai and Quoc V. Le. 2015. [Semi-supervised sequence learning](#). In *Advances in Neural Information Processing Systems 28*, pages 3079–3087.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Qing Dou and Kevin Knight. 2012. [Large scale decipherment for out-of-domain machine translation](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 266–275, Jeju Island, Korea. Association for Computational Linguistics.
- Qing Dou and Kevin Knight. 2013. [Dependency-based decipherment for resource-limited machine translation](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1668–1676, Seattle, Washington, USA. Association for Computational Linguistics.
- Long Duong, Hiroshi Kanayama, Tengfei Ma, Steven Bird, and Trevor Cohn. 2017. [Multilingual training of crosslingual word embeddings](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 894–904, Valencia, Spain. Association for Computational Linguistics.
- Umberto Eco and James Fentress. 1995. *The search for the perfect language*. Blackwell Oxford.
- Manaal Faruqui and Chris Dyer. 2014. [Improving vector space word representations using multilingual correlation](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 462–471, Gothenburg, Sweden. Association for Computational Linguistics.
- Goran Glavaš, Robert Litschko, Sebastian Ruder, and Ivan Vulić. 2019. [How to \(properly\) evaluate crosslingual word embeddings: On strong baselines, comparative analyses, and some misconceptions](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 710–721, Florence, Italy. Association for Computational Linguistics.
- Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. [BiBOWA: Fast bilingual distributed representations without word alignments](#). In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 748–756, Lille, France. PMLR.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. [Learning word vectors for 157 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Edouard Grave, Armand Joulin, and Quentin Berthet. 2019. [Unsupervised alignment of embeddings with wasserstein procrustes](#). In *Proceedings of Machine Learning Research*, volume 89, pages 1880–1890. PMLR.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. [The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English](#). 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 6097–6110, Hong Kong, China. Association for Computational Linguistics.
- Michael Hahn and Marco Baroni. 2019. [Tabula nearly rasa: Probing the linguistic knowledge of character-level neural language models trained on unsegmented text](#). *Transactions of the Association for Computational Linguistics*, 7:467–484.
- Geert Heyman, Bregt Verreet, Ivan Vulić, and Marie-Francine Moens. 2019. [Learning unsupervised multilingual word embeddings with incremental multilingual hubs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1890–1902, Minneapolis, Minnesota. Association for Computational Linguistics.
- Geert Heyman, Ivan Vulić, and Marie-Francine Moens. 2017. [Bilingual lexicon induction by learning to combine word-level and character-level representations](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1085–1095, Valencia, Spain. Association for Computational Linguistics.
- Yedid Hoshen and Lior Wolf. 2018. [Non-adversarial unsupervised word translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 469–478, Brussels, Belgium. Association for Computational Linguistics.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Tsung-Yuan Hsu, Chi-Liang Liu, and Hung-yi Lee. 2019. [Zero-shot reading comprehension by crosslingual transfer learning with multi-lingual language representation model](#). 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 5933–5940, Hong Kong, China. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. **XTREME: A Massively Multilingual Multi-task Benchmark for Evaluating Cross-lingual Generalization**. *arXiv preprint arXiv:2003.11080*.
- David Kamholz, Jonathan Pool, and Susan Colowick. 2014. **PanLex: Building a resource for pan-lingual lexical translation**. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3145–3150, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Yova Kementchedjheva, 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 3327–3332, Hong Kong, China. Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. **Unsupervised machine translation using monolingual corpora only**. In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018b. **Phrase-based & neural unsupervised machine translation**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5039–5049, Brussels, Belgium. Association for Computational Linguistics.
- Patrick Lewis, Barlas Öğuz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2019. **MLQA: Evaluating Cross-lingual Extractive Question Answering**. *arXiv preprint arXiv:1910.07475*.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Bruce Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Rangan Majumder, and Ming Zhou. 2020. **Xglue: A new benchmark dataset for cross-lingual pre-training, understanding and generation**. *arXiv preprint arXiv:2004.01401*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. **Multilingual denoising pre-training for neural machine translation**. *arXiv preprint arXiv:2001.08210*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. **Bilingual word representations with monolingual quality in mind**. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 151–159, Denver, Colorado. Association for Computational Linguistics.
- Kelly Marchisio, Kevin Duh, and Philipp Koehn. 2020. **When does unsupervised machine translation work?** *arXiv preprint arXiv:2004.05516*.
- Benjamin Marie and Atsushi Fujita. 2018. **Unsupervised neural machine translation initialized by unsupervised statistical machine translation**. *arXiv preprint arXiv:1810.12703*.
- Benjamin Marie, Haipeng Sun, Rui Wang, Kehai Chen, Atsushi Fujita, Masao Utiyama, and Eiichiro Sumita. 2019. **NICT’s unsupervised neural and statistical machine translation systems for the WMT19 news translation task**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 294–301, Florence, Italy. Association for Computational Linguistics.
- Thomas Mayer and Michael Cysouw. 2014. **Creating a massively parallel Bible corpus**. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3158–3163, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. **Learned in translation: Contextualized word vectors**. In *Advances in Neural Information Processing Systems 30*, pages 6294–6305.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013a. **Exploiting similarities among languages for machine translation**. *arXiv preprint arXiv:1309.4168*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013b. **Distributed representations of words and phrases and their compositionality**. In *Advances in Neural Information Processing Systems 26*, pages 3111–3119.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. **GloVe: Global vectors for word representation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. **Deep contextualized word representations**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.

- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. [Massively multilingual transfer for NER](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 151–164, Florence, Italy. Association for Computational Linguistics.
- Sujith Ravi and Kevin Knight. 2011. [Deciphering foreign language](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 12–21, Portland, Oregon, USA. Association for Computational Linguistics.
- Shuo Ren, Zhirui Zhang, Shujie Liu, Ming Zhou, and Shuai Ma. 2019. Unsupervised neural machine translation with SMT as posterior regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 241–248.
- 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)*, pages 390–394, Melbourne, Australia. Association for Computational Linguistics.
- 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, Brussels, Belgium. 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.
- Tal Schuster, Ori Ram, Regina Barzilay, and Amir Globerson. 2019. [Cross-lingual alignment of contextual word embeddings, with applications to zero-shot dependency parsing](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1599–1613, Minneapolis, Minnesota. Association for Computational Linguistics.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2019a. [WikiMatrix: Mining 135M Parallel Sentences](#). *arXiv preprint arXiv:1907.05791*.
- Holger Schwenk and Xian Li. 2018. [A corpus for multilingual document classification in eight languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, and Armand Joulin. 2019b. [CC-Matrix: Mining Billions of High-Quality Parallel Sentences on the WEB](#). *arXiv preprint arXiv:1911.04944*.
- Aditya Siddhant, Melvin Johnson, Henry Tsai, Naveen Arivazhagan, Jason Riesa, Ankur Bapna, Orhan Firat, and Karthik Raman. 2019. [Evaluating the Cross-Lingual Effectiveness of Massively Multilingual Neural Machine Translation](#). *arXiv preprint arXiv:1909.00437*.
- Samuel L. Smith, David H. P. Turban, Steven Hamblin, and Nils Y. Hammerla. 2017. [Offline bilingual word vectors, orthogonal transformations and the inverted softmax](#). In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*.
- Benjamin Snyder, Tahira Naseem, Jacob Eisenstein, and Regina Barzilay. 2008. [Unsupervised multilingual learning for POS tagging](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 1041–1050, Honolulu, Hawaii. 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, Melbourne, Australia. Association for Computational Linguistics.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tiejun Liu. 2019. [MASS: Masked sequence to sequence pre-training for language generation](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 5926–5936, Long Beach, California, USA. PMLR.
- Dario Stojanovski, Viktor Hangya, Matthias Huck, and Alexander Fraser. 2019. [The LMU munich unsupervised machine translation system for WMT19](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 393–399, Florence, Italy. Association for Computational Linguistics.
- Clara Vania and Adam Lopez. 2017. [From characters to words to in between: Do we capture morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2016–2027, Vancouver, Canada. Association for Computational Linguistics.
- Ivan Vulić, Goran Glavaš, Roi Reichart, and Anna Korhonen. 2019. [Do we really need fully unsupervised cross-lingual embeddings?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International*

Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 4407–4418, Hong Kong, China. Association for Computational Linguistics.

Ivan Vulić and Anna Korhonen. 2016. [On the role of seed lexicons in learning bilingual word embeddings](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 247–257, Berlin, Germany. Association for Computational Linguistics.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Shijie Wu, Alexis Conneau, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Emerging cross-lingual structure in pretrained language models](#). *arXiv preprint arXiv:1911.01464*.

Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). 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 833–844, Hong Kong, China. Association for Computational Linguistics.

Ruochen Xu, Yiming Yang, Naoki Otani, and Yuexin Wu. 2018. [Unsupervised cross-lingual transfer of word embedding spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2465–2474, Brussels, Belgium. Association for Computational Linguistics.

Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. 2017a. [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, Vancouver, Canada. Association for Computational Linguistics.

Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. 2017b. [Earth mover’s distance minimization for unsupervised bilingual lexicon induction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1934–1945, Copenhagen, Denmark. Association for Computational Linguistics.