

XeroAlign: Zero-Shot Cross-lingual Transformer Alignment

Milan Gritta and Ignacio Iacobacci

Huawei Noah's Ark Lab, London, UK

{milan.gritta, ignacio.iacobacci}@huawei.com

Abstract

The introduction of transformer-based cross-lingual language models brought decisive improvements to multilingual NLP tasks. However, the lack of labelled data has necessitated a variety of methods that aim to close the gap to high-resource languages. Zero-shot methods in particular, often use translated task data as a training signal to bridge the performance gap between the source and target language(s). We introduce **XeroAlign**, a simple method for task-specific alignment of cross-lingual pretrained transformers such as XLM-R. XeroAlign uses translated task data to encourage the model to generate similar sentence embeddings for different languages. The XeroAligned XLM-R, called XLM-RA, shows strong improvements over the baseline models to achieve state-of-the-art zero-shot results on three multilingual natural language understanding tasks. XLM-RA performs on par with state-of-the-art models on a cross-lingual adversarial paraphrasing task and its text classification accuracy exceeds that of XLM-R trained with labelled data.

1 Introduction

Recently, large transformer-based (Vaswani et al., 2017) pretrained language models have achieved state-of-the-art (SOTA) performance on several monolingual and multilingual NLP tasks (Wang et al., 2019a; Hu et al., 2020b). Transfer learning enabled the self-supervised pretraining on unlabelled datasets to learn linguistic features such as syntax and semantics in order to improve tasks with limited training data (Wang et al., 2019b). Pretrained cross-lingual language models (PXLMS) have soon followed to learn general linguistic features and properties of dozens of languages (Lample and Conneau, 2019; Xue et al., 2020). For multilingual tasks in particular, adequate labelled data is usually only available for a few well-resourced languages such as English. Zero-shot approaches

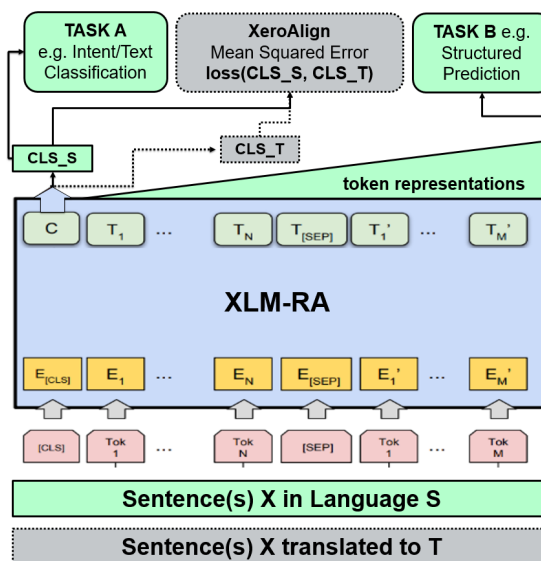


Figure 1: The XeroAligned XLM-R model (XLM-RA) for cross-lingual NLU. XeroAlign is added to the training loop to encourage embeddings in different languages to be similar, enabling a zero-shot classifier transfer.

were introduced to transfer the knowledge to languages without the requisite training data. To this end, we introduce **XeroAlign**, a conceptually simple and efficient method for task-specific alignment of sentence embeddings generated by PXLMS, aimed at effective zero-shot cross-lingual transfer. XeroAlign is an auxiliary training objective, which uses translated data to bring the performance in the target language closer to the source (labelled) language, as illustrated in Figure 1. We apply our proposed method to the publicly available XLM-R transformer (Conneau et al., 2020) but instead of pursuing large-scale model alignment with general parallel corpora such as Europarl (Koehn, 2005), we show that a simplified, task-specific model alignment is an effective and efficient approach to zero-shot transfer for cross-lingual natural language understanding (XNLU).

We evaluate XeroAlign on four datasets (two different tasks) that cover 11 unique languages. The XeroAligned XLM-R model (XLM-RA) achieves SOTA scores on three XNLU datasets, exceeds the text classification performance of XLM-R trained with labelled data and performs on par with SOTA models on an adversarial paraphrasing task.

2 Background

We first outline the necessary background on Contrastive Learning (Becker and Hinton, 1992) and transformer-based PXLMS. Contrastive learning (CL) is a framework designed for self-supervised representation learning. Recent examples include Momentum Contrast (MoCo) (He et al., 2020) and SimCLR (Chen et al., 2020), both of which achieved strong improvements on image classification. The essence of CL is to generate representations that are similar for positive examples and dissimilar for negative examples. CL-based methods in cross-lingual NLP replace negative samples (formerly augmented images) with random sentences in the target language, typically thousands of them. Positive examples comprise sentences translated into the target language. While CL is applicable to task-agnostic model alignment, large batches of random negative samples are infeasible for small datasets. Negative samples drawn randomly from a small dataset can result in duplicates (false negatives), which is why our alignment uses only positive samples. The CL-based alignments in section 2.1 aim to improve generic cross-lingual representations with large (parallel) datasets. In contrast, we align the PXLMS only with translated task data, making our approach simpler and more efficient, showing a strong zero-shot transfer.

Transformer-based PXLMS For transformer-based PXLMS, two basic types of representations are commonly used: 1) A sentence embedding for tasks such as text classification (Conneau et al., 2018) or sentence retrieval (Zweigenbaum et al., 2018), which use the [CLS] representation of the full input sequence, and 2) Token embeddings, which are used for structured prediction (Pan et al., 2017) or Q&A (Lewis et al., 2020), requiring each token’s contextualised representation for a per-token inference. While our method uses the [CLS] embedding, other approaches based on Contrastive Learning have used both types of representations to obtain a sentence embedding.

2.1 Related Work

In order to cluster prior work, we formulate a taxonomy (Table 1) for the purposes of positioning our method in the appropriate context. The relevant zero-shot transfer approaches can be grouped by a) whether the alignment is targeted at each task, i.e. is task-specific [TS] or is task-agnostic [TA] and b) whether the alignment is applied to the model [MA] or data [DA]. Our paper falls mostly into the [MA,TS] category although methodological similarities are also found in the [MA,TA] group.

Groups	<i>Task-Specific</i>	<i>Task-Agnostic</i>
<i>Data Align</i>	[DA,TS]	No relevant work
<i>Model Align</i>	[MA,TS]	[MA,TA]

Table 1: An approximate taxonomy of prior work.

[MA,TA] Hu et al. (2020a) have proposed two objectives for cross-lingual zero-shot transfer a) sentence alignment and b) word alignment. While CL is not mentioned, the proposed sentence alignment closely resembles contrastive learning with one encoder (e.g. SimCLR). Taking the average of the contextualised token representations as the input representation (as an alternative to the [CLS] token), the model predicts the correct translation of the sentence within a batch of negative samples. An improvement is observed for text classification tasks and sentence retrieval but not structured prediction. The alignment was applied to a 12-layer multilingual BERT and the scores are comparable to the translate-train baseline (translate data and train normally). Instead, we use one of the best publicly available models, XLM-R from Huggingface, as our starting point since an improvement in a weaker baseline is not guaranteed to work in a stronger model that may have already subsumed those upgrades. Contrastive alignment based on MoCo with two PXLMS encoders was proposed by Pan et al. (2020). Using an L2 normalised [CLS] token with a non-linear projection as the input representation, the model was aligned on 250K to 2M parallel sentences with added Translation Language Modelling (TLM) and a code-switching augmentation. No ablation for MoCo was provided to estimate its effect although the combination of all methods did provide improvements with multilingual BERT as the base learner. Another model inspired by CL is InfoXLM (Chi et al., 2020). InfoXLM is pretrained with TLM, multilingual

Masked Language Modelling and Cross-lingual Contrastive Learning called XLCo. Like MoCo, they employ two encoders that use the [CLS] token (or the uppermost layer average) as the sentence representation, taken from layers 8 (base model) and 12 (large model). Ablation showed a 0.2-0.3 improvement in accuracy for XNLI and MLQA (Lewis et al., 2020). Reminiscent of earlier work on cross-lingual alignment (Hermann and Blunsom, 2014), the task-agnostic sentence embedding model (Feng et al., 2020) called LaBSe (Language-agnostic BERT sentence embeddings) uses the [CLS] representations of two BERT encoders (compared to our single encoder) with a margin loss and 6 billion parallel sentences to generate multilingual representations. While similarities exist, our multi-task alignment is an independently devised, more efficient, task-specific and a simplified version of the aforementioned approaches.

[DA,TS] Zero-shot cross-lingual models often make use of machine translation to provide a training signal. This is an uncomplicated transformation for text classification tasks given that adequate machine translation models exist for many language pairs. However, for structured prediction tasks such as Slot Filling or Named Entity Recognition, the non-trivial task of *aligning token labels* can lead to an improved cross-lingual transfer. One of the most commonly used word alignment methods is fastalign (Dyer et al., 2013). Frequently used as a baseline, it involves aligning the word indices in parallel sentences in an unsupervised manner, prior to regular supervised learning. In some scenarios, fastalign can approach SOTA scores for slot filling (Schuster et al., 2019), however, the quality of alignment varies between languages and can even degrade performance (Li et al., 2021) below baseline. An alternative data alignment approach called CoSDA (Qin et al., 2020) uses code-switching as data augmentation. Random words in the input are translated and replaced to make model training highly multilingual, leading to improved cross-lingual transfer. Attempts were also made to automatically learn how to code-switch (Liu et al., 2020). While improvements were reported, it’s uncertain how much SOTA models would benefit.

[MA,TS] Continuing with label alignment for slot filling, Xu et al. (2020) tried to predict and align slot labels jointly during training instead of modifying data labels explicitly before fine-tuning. While soft-align improves on fastalign, the diffi-

culty of label alignment makes it challenging to improve on the SOTA. For text classification tasks such as Cross-lingual Natural Language Inference (Conneau et al., 2018), an adversarial cross-lingual alignment was proposed by Qi and Du (2020). Adding a self-attention layer on top of multilingual BERT (Devlin et al., 2019) or XLM (Lample and Conneau, 2019), the model learns the XNLI task while trying to fool the language discriminator in order to produce language-agnostic input representations. While improvements over baselines were reported, the best scores were around 2-3 points behind the standard XLM-R model.

3 Methodology

We introduce **XeroAlign**, a conceptually simple and efficient method for task-specific alignment of sentence embeddings generated by PXLMS, aimed at effective zero-shot cross-lingual transfer. XeroAlign is an auxiliary training objective that is jointly optimised with the primary task, e.g. text classification and/or slot filling, as shown in Figure 1. We use standard architecture for each task and only add the minimum required number of new parameters. For text classification tasks, we use the [CLS] token of the PXML as our pooled sentence representation. A linear classifier (hidden size \times number of classes) is learnt on top of the [CLS] embedding using cross-entropy as the loss function (TASK A in Figure 1). For slot filling, we use the contextualised representations of each token in the input sequence. Once again, a linear classifier (hidden size \times number of slots) is learnt with a cross-entropy loss (TASK B in Figure 1).

Algorithm 1 shows a standard training routine augmented with XeroAlign. Let $PXML$ be a pre-trained cross-lingual transformer language model, X be the standard English training data and U be the machine translated utterances (from X). The English utterances were translated into each target language using our internal machine translation service. A public online translator e.g. Google Translate can also be used. For the PAWS-X task, we use the public version of the translated data¹. We then obtain the CLS_S and CLS_T embeddings by taking the first token of the $PXML$ output sequence for the source x_s and target x_t sentences respectively. Using a Mean Squared Error loss

¹<https://github.com/google-research-datasets/paws>

Algorithm 1 The XeroAlign algorithm.

```
1:  $PXLM \leftarrow$  Pretrained Cross-lingual LM
2:  $X \leftarrow$  Training data tuples in English
3:  $U \leftarrow$  Utterances translated into Target Lang.
4:  $sim \leftarrow$  similarity function e.g. MSE

5: # training loop
6: for  $(x_s, y), x_t \in X, U$  do
7:    $task\_loss \leftarrow task\_loss\_fn(x_s, y)$ 
8:    $CLS_S \leftarrow PXLM(x_s)$ 
9:    $CLS_T \leftarrow PXLM(x_t)$ 
10:   $align\_loss \leftarrow sim(CLS_S, CLS_T)$ 
11:   $total\_loss \leftarrow task\_loss + align\_loss$ 
12:  # update model parameters
13: end for
```

function as our similarity function sim , we compute the distance/loss between CLS_S and CLS_T . The sum of the losses ($total_loss$) is then backpropagated normally. We have conducted all XeroAlign training as multi-task learning for the following reason. When the $PXLM$ is aligned first, followed by primary task training, the $PXLM$ exhibits poor zero-shot performance. Similarly, learning the primary task first, followed by XeroAlign fails as the primary task is partially unlearned during alignment. This is most likely due to the catastrophic forgetting problem (Goodfellow et al., 2013).

3.1 Experimental Setup

In order to make our method easily accessible and reproducible², we use the publicly available XLM-R transformer from Huggingface (Wolf et al., 2019) built on top of PyTorch (Paszke et al., 2019). We set a single seed for all experiments and a similar learning rate for each dataset. No hyperparameter sweep was conducted to ensure a robust, low-resource, real-world deployment and to make a fair comparison with SOTA models. XLM-R was XeroAligned over 10 epochs and optimised using Adam (Kingma and Ba, 2015) and a OneCycleLR (Smith and Topin, 2019) scheduler.

3.2 Datasets

We evaluate XeroAlign with four datasets covering 11 unique languages (en, de, es, fr, th, hi, ja, ko, zh, tr, pt) across two tasks (intent classification + slot filling, paraphrase detection).

²Resources are available by email (Milan Gritta) or download from <https://www.aclweb.org/anthology/> or <https://github.com/huawei-noah/xnlu>.

PAWS-X (Yang et al., 2019) is a multilingual version of PAWS (Zhang et al., 2019), a binary classification task for identifying paraphrases. Examples were sourced from Quora Question Pairs³ and Wikipedia, chosen to mislead simple ‘word overlap’ models. PAWS-X contains random examples drawn from PAWS (just under 4K for the development and test set) covering seven languages (en, de, es, fr, ja, ko, zh) and totalling almost 28,000 human translated paraphrases. Each train set contains approximately 49.5K machine translated examples.

MTOD is a Multilingual Task-Oriented Dataset provided by Schuster et al. (2019). It covers three domains (alarm, weather, reminder) and three languages of different sizes: English (43K), human-translated Spanish (8.3K) and Thai (5K). MTOD comprises two correlated NLU tasks, intent classification and slot filling. The SOTA scores are reported by Li et al. (2021) and Schuster et al. (2019).

MTOP is a Multilingual Task-Oriented Parsing dataset provided by Li et al. (2021) that covers interactions with a personal assistant (intent recognition and slot filling tasks). We use the standard flat version, which has the SOTA reported by Li et al. (2021). A compositional version of the data designed for nested queries is also provided. MTOP contains 100K+ human translated examples in 6 languages (en, de, es, fr, th, hi) and 11 domains.

MultiATIS++ by Xu et al. (2020) is an extension of the multilingual version of ATIS (Upadhyay et al., 2018), initially translated into Hindi and Turkish. Six new human-translated⁴ languages (de, es, fr, zh, ja, pt) were added with ~ 4 times as many examples each (around 6K per language) for 9 languages. These datasets are based on the English-only ATIS (Price, 1990) featuring users interacting with an information service. Once again, intent recognition and slot filling are evaluated.

3.3 Metrics

We use standard evaluation metrics, that is, accuracy for paraphrase detection and intent classification, F-Score⁵ for slot filling.

³<https://www.quora.com/q/quoradata/First-Quora-Dataset-Release-Question-Pairs>

⁴We have encountered some minor issues with slot annotations. Around 60-70 entities across 5 languages (fr, zh, hi, ja, pt) had to be corrected as the number of slot tags did not agree with the number of tokens in the sentence. However, this only concerns a tiny fraction of the $\sim 400k+$ tokens covered.

⁵<https://pypi.org/project/segeval/>

Model	Spanish	French	German	Hindi	Thai	Average
XLM-R Target	95.9 / 91.2	95.5 / 89.6	96.6 / 88.3	95.1 / 89.1	94.8 / 87.7	95.6 / 89.2
XLM-R 0-shot	91.9 / 84.3	93.0 / 83.7	87.5 / 80.7	91.4 / 76.5	87.6 / 55.6	90.3 / 76.2
XLM-RA	96.6 / 84.4	96.5 / 83.3	95.7 / 84.5	95.2 / 80.1	94.1 / 69.1	95.6 / 80.3
Li et al. (2021)	96.3 / 84.8	95.1 / 82.5	94.8 / 83.1	94.2 / 76.5	92.1 / 65.6	94.5 / 77.9

Table 2: MTOP results as Intent Classification Accuracy / Slot Filling F-Score. Best English score: 97.3 / 93.9.

4 Results and Analysis

We use ‘XLM-R Target’ to refer to models trained on each labelled target language. We also provide zero-shot scores denoted ‘XLM-R 0-shot’, the ‘XLM-RA’ results and the reported SOTA figures. For PAWS-X, we provide a second baseline called ‘Translate-Train’, which comprises the union of Target and English train data. Scores are given for the large⁶ model unless specified otherwise.

The XeroAligned XLM-R attains SOTA scores on 3 XNLU datasets. The intent classification (+1.1) and slot filling (+2.4) scores, averaged over 5 MTOP languages and shown in Table 2, improved on an XLM-R model with translated utterances, slot label projection and distant supervision (Li et al., 2021). For the MultiATIS++ dataset (Table 5), XLM-RA shows an improved intent accuracy (+1.1) and F-Score (+3.2) over 8 languages, compared to a large multilingual BERT with translated utterances and slot label softalign (Xu et al., 2020). The MTOD classification accuracy (+1.3) and slot tagging F-Score (+5.0) have improved over an XLM-R with translated utterances, slot label projection and distant supervision (Li et al., 2021). MTOD is the only dataset where the base model outperforms (albeit marginally) the large transformer, as seen in Table 6. Finally, we compare our intent classification accuracy (+8.1) and slot filling F-Score (+8.7) for MTOD to a BiLSTM with translated utterances and slot label projection (Schuster et al., 2019), which had the SOTA F-Score for Thai.

On the adversarial paraphrase task (PAWS-X, Table 3), averaged over 7 languages, XLM-RA scores marginally higher (+0.1 accuracy) than VECO (Luo et al., 2020), a variable cross-lingual encoder-decoder and marginally lower (-0.2 accuracy) than FILTER (Fang et al., 2020), an enhanced cross-

lingual fusion model, which was the SOTA until 01/2021. We now turn our attention to the improvements over the XLM-R 0-shot model.

4.1 Text Classification

The intent classification accuracy of XLM-RA exceeds that of an XLM-R trained with labelled data, averaged across three task-oriented XNLU datasets (Tables 2, 5, 6). Starting from a competitive baseline model, XeroAlign improves intent accuracy by ~5-10 points (similar for the base model, see Table 7 in Section 4.4). The benefits of cross-lingual alignment are particularly evident in low-resource languages (tr, hi, th), which is encouraging for applications with limited resources. Paraphrase detection is another instance of text classification. We report average XLM-RA accuracy in Table 3, which exceeds both Target and the Translate-Train scores by over 1 point and by almost 3 points over XLM-R 0-shot (even larger improvements for the base model).

Note that the amount of labelled data is comparable for XeroAlign and Target (except Thai, Spanish in MTOD and Turkish, Hindi in MultiATIS++) hence there is no more than a negligible advantage from the additional data. As the English scores show the highest performance (~1.5 points higher on average) compared to the Target languages, we hypothesise that transferring this advantage from a high-resource language via XeroAlign is the primary reason behind its effectiveness compared to using target data directly. Given that Target performance has recently been exceeded with MoCo (He et al., 2020) and the similarities between contrastive learning and XeroAlign, our finding seems in line with recent work, which is subject to ongoing research (Zhao et al., 2020).

4.2 Structured Prediction

While XLM-RA is able to exceed Target accuracy for text classification tasks, even our best F-Scores

⁶Large=24 layers, 550M par, Base=12 layers, 270M par.

Model	EN	DE	ES	FR	JA	KO	ZH	Average
XLM-R Target	95.6	90.9	92.5	92.4	85.1	86.4	87.2	90.0
XLM-R Translate-Train	95.7	91.6	92.3	92.5	85.2	85.8	87.7	90.1
XLM-R 0-shot	95.6	91.0	91.1	91.9	81.7	81.6	85.4	88.3
Luo et al. (2020)	96.4	93.0	93.0	93.5	87.2	86.8	87.9	91.1
XLM-RA	95.8	92.9	93.0	93.9	87.1	87.1	88.9	91.2
Fang et al. (2020)	95.9	92.8	93.0	93.7	87.4	87.6	89.6	91.4

Section 4.5 experiment below: aligning with development/test set utterances but no task labels.

XLM-RA (Exp)	95.8	94.2	94.4	94.8	91.6	92.6	92.1	93.6
--------------	------	------	------	------	------	------	------	------

Table 3: PAWS-X results as Paraphrase Classification Accuracy.

for slot filling are between 8 and 19 points behind the Target model. This is despite a strong average improvement of +4.1 on MTOP, +5.7 on Multi-ATIS++ and +5.2 on MTOD (a greater benefit for the base model). We think that the gap is primarily down to the difficulty of the sequence labelling task. That is, zero-shot text classification is ‘easier’ than zero-shot slot filling, which is manifested by a ~ 10 -20 point gap between scores. Sentences in various languages have markedly different input lengths and token order thus word-level inference in cross-lingual settings becomes significantly more challenging than sentence-level prediction because syntax plays a less critical role in sequence classification.

Another reason, related to XeroAlign’s architecture and training, is our choice to align the PXML on the [CLS] embedding, which is subsequently used ‘as is’ for text classification tasks. Aligning individual token representations through the [CLS] embedding improves structured prediction as well, however, as the token embeddings are not directly used, the Multi-Head Attention parameters in the uppermost transformer layer never receive any gradient updates from XeroAlign. Closing this gap is a challenging opportunity, which we reserve for future work. Once again, the languages with lower NLP resources (th, hi, tr) tend to benefit the most from cross-lingual alignment.

4.3 XeroAlign Generalisation

In this subsection, we briefly investigate the generalisation of XeroAlign to unseen languages, taking the PAWS-X task as our use case. We are interested in finding out whether aligning on *just one language* has any zero-shot benefits for other lan-

guages. Table 4 shows the XLM-RA results when aligned on a single language (rows) and tested on other languages (columns).

-	EN	DE	ES	FR	JA	KO	ZH	AVE
DE	96.0	92.9	92.3	92.6	84.0	84.5	86.5	89.8
ES	95.9	92.6	93.0	93.1	83.9	84.1	86.4	89.9
FR	95.9	92.5	92.9	93.9	83.9	84.1	86.9	90.0
JA	96.0	92.6	91.8	93.1	87.1	87.4	87.9	90.8
KO	95.7	92.6	92.0	92.7	80.6	87.1	87.3	90.5
ZH	95.5	92.0	92.6	92.7	86.3	86.2	88.9	90.6
EU	96.2	92.5	93.0	94.1	84.9	85.2	87.1	90.4
AS	96.0	93.0	92.1	92.7	85.9	87.6	88.4	90.8

Table 4: XLM-RA aligned on one PAWS-X language (rows), evaluated on others (columns). AVE = average. EU = European languages, AS = Asian languages.

We can see that aligning on Asian languages (Japanese in particular) attains the best average improvement compared to aligning with European languages. This seems to reflect the known performance bias of XLM-R towards (high-resource) European languages, all of which show a strong improvement, regardless of language. Aligning only on European languages (de, es, fr) improves the average to 90.4 but aligning on Asian languages (zh, ko, ja) does not improve over Japanese (90.8). In any case, it is notable that the XLM-R model XeroAligned on *just a single language* is able to carry this advantage well beyond a single language thus improve average accuracy by 1.5-2.5 points over baseline (88.3) from Table 3. This effect is even stronger for MTOP (+4 accuracy, +3 F-Score).

4.4 Smaller Language Models

The XeroAligned XLM-R base model shows an even greater relative improvement than its larger counterpart with 24 layers and 550M parameters. To this end, we report the results for the XLM-

Model	DE	ES	FR	TR	HI	ZH	PT	JA	AVE
XLM-R Target	97.0/95.3	97.3/87.9	97.8/93.8	80.6/74.0	89.7/84.1	95.5/95.9	97.2/94.1	95.5/92.6	93.8/89.7
XLM-R 0-shot	96.4/84.8	97.0/85.5	95.3/81.8	76.2/41.2	91.9/68.2	94.3/82.5	90.9/81.9	89.8/77.6	91.5/75.5
XLM-RA	97.6/84.9	97.8/85.9	95.4/ 81.4	93.4/ 70.6	94.0/79.7	96.4/83.3	97.6/79.9	96.1/83.5	96.0/81.2
Jain et al. (2019)	96.0/87.5	97.0/84.0	97.0/79.8	93.7/44.8	92.4/77.2	95.2/ 85.1	96.5/ 81.7	88.5/82.6	94.5/77.8
Xu et al. (2020)	96.7/ 89.0	97.2/76.4	97.5/79.6	93.7/61.7	92.8/78.6	96.0/83.3	96.8/76.3	88.3/79.1	94.9/78.0

Table 5: MultiATIS++ Intent Classification Accuracy / Slot Filling F1. English model: 97.9/97. AVE = average.

Model	Spanish	Thai	AVE
§ Target (B)	98.7/89.1	96.8/93.1	97.8/91.1
§ Target (L)	98.8/89.8	97.8/94.4	98.3/92.1
§ 0-shot (B)	90.7/70.1	71.9/53.1	81.3/61.6
§ 0-shot (L)	97.1/85.7	82.8/47.7	90.0/66.7
XLM-RA (B)	98.9/86.9	97.9/ 60.2	98.4/ 73.6
XLM-RA (L)	99.2/88.4	98.4/57.3	98.8/72.9
Schuster et al.	85.4/72.9	95.9/55.4	90.7/64.2
Li et al.	98.0/83.0	96.9/52.8	97.5/67.9

Table 6: MTOOD results as Intent Classification Accuracy / Slot Filling F-Score. Our best English score: 99.3/96.6, (B) = Base, (L) = Large, § = XLM-R model.

RA base model (12 layers, 270M parameters) in Table 7 as the average scores over all languages for MTOP, PAWS-X, MTOOD and MultiATIS++. We use a relative % improvement over XLM-R 0-shot to compare the models fairly. We observe that the paraphrase detection accuracy improves by 3.3% for the large PXML versus 6.5% for the base model. Across three XNLU datasets, the large XeroAligned XLM-R improves over the XLM-R 0-shot by 9.5% versus 14.2% for the base model on slot filling and by 7.1% (large) versus 19.8% (base) on text/intent classification.

Model	MTOP	PAWS-X	M-ATIS	MTOOD
§ Target	94.0/88.1	85.2	89.0/86.3	97.6/92.2
§ 0-shot	80.8/68.9	81.7	76.9/65.0	80.1/64.8
XLM-RA	93.3/78.9	87.0	93.0/73.4	98.5/74.7

Table 7: The XLM-R(A) base model averages as intent classification accuracy / slot filling F-Score (or paraphrase accuracy for PAWS-X). § = XLM-R model.

This lets applications with lower computational budgets achieve competitive performance with our efficient cross-lingual alignment method for transformed-based PXMLs. In fact, the XLM-RA base model is able to reach up to 90-95% of the performance of its larger sibling using significantly lower computational resources.

4.5 Discussion

XLM-RA’s intent classification accuracy is on average within ~ 1.5 points of English accuracy across three XNLU datasets. However, the PAWS-X paraphrase detection accuracy is almost 5 points below English models, which is the case for other state-of-the-art PXMLs in Table 3. Why does the XLM-R struggle to generalise more on this task for languages other than English? A combination of several factors is likely at play. Minor translation errors may account for some of this deficit since the training is done using the publicly available machine-translated data. XeroAlign also leverages this data to align the multilingual representations of the PXML. It is the nature of the dataset makes translation errors more likely for PAWS-X than for the task-oriented XNLU as the differences between paraphrased sentences are very subtle and challenging to discern (and to translate accurately). With that said, we think that the greater than expected deficit may be primarily caused by 1) a domain/topic shift within the dataset and 2) a possible data leakage for English. The original PAWS data (Zhang et al., 2019) was sourced from Quora Question Pairs and Wikipedia with neither being limited to any particular domain. As the English Wikipedia provides a large portion (relative to other languages) of the training data for XLM-R, it is possible that some of the English PAWS sentences may have been seen in training, which could explain the smaller generalisation gap for English.

In order to find out whether the accuracy gap will diminish, we artificially remove the hypothetical domain shift by using the parallel utterances (but not task labels) from the development and test sets. We thus XeroAlign the XLM-R on an extended vocabulary that may not be present in the train set. We observe that the (Exp) model in Table 3 shows an average improvement of over 2 points compared to the best XLM-RA and other SOTA models. This suggests that the increased generalisation gap may be caused by a domain shift for non-English lan-

guages on this task. When that topic shift gets (perhaps artificially) removed, the model is able to bring accuracy back within ~ 2 points of the English model (in line with XNLU tasks). Note that this gap can be masked for English due to the language biases in data used for pretraining.

In section 2, we outlined the most conceptually similar methods that conducted large-scale model pretraining with task-agnostic sentence alignment as part of the training routine (Hu et al., 2020a; Feng et al., 2020; Pan et al., 2020; Chi et al., 2020). Where ablation studies were provided, the average improvement attributed to contrastive alignment was ~ 0.2 - 0.3 points (though the tasks were slightly different). While we do not directly compare XeroAlign to contrastive alignment, it seems that a task-specific alignment may be a more effective and efficient technique to improve zero-shot transfer, given the magnitude of our results. This leads us to conclude that the effectiveness of our method comes primarily from cross-lingual alignment of the task-specific vocabulary. Language is inherently ambiguous, the semantics of words and phrases shift somewhat from topic to topic, therefore, a cross-lingual alignment of sentence embeddings *within the context of the target task* should lead to better results. Our simplified, lightweight method only uses translated task utterances, a single encoder model and positive samples, the alignment of which is challenging enough without arbitrary negative samples. In fact, this is the main barrier for applying contrastive alignment in task-specific NLP scenarios, i.e. the lack of carefully constructed negative samples. For smaller datasets, random negative samples would mean that the task is either too easy to solve, resulting in no meaningful learning or the model would receive conflicting signals by training on false positive examples, leading to degenerate learning.

4.6 Future Work

Our recommendations for avenues of promising follow-up research involve any of the following: i) aligning additional tasks such as Q&A, Natural Language Inference, Sentence Retrieval, etc. ii) including additional languages, especially low-resource ones (Joshi et al., 2020) and iii) evaluating large-scale, task-agnostic alignment of PXLMS followed by task-specific alignment, which is reminiscent of the common transfer learning paradigm of pretraining with Masked Language Modelling

before fine-tuning on the target task. To that end, there is already some emergent work on monolingual fine-tuning with an additional contrastive loss (Gunel et al., 2020). For the purposes of multilingual benchmarks (Hu et al., 2020b; Liang et al., 2020) or other pure empirical pursuits, an architecture or a language-specific hyperparameter search should optimise XLM-RA for significantly higher performance as the large transformer does not always outperform its smaller counterpart and because our hyperparameters remained fixed for all languages. Most importantly, the follow-up work needs to improve zero-shot transfer for cross-lingual *structured prediction* such as Named Entity Recognition (Pan et al., 2017), POS Tagging (Nivre et al., 2016) or Slot Filling (Schuster et al., 2019), which is still lagging behind Target scores.

5 Conclusions

We have introduced **XeroAlign**, a conceptually simple and efficient method for task-specific alignment of multilingual sentence embeddings generated by PXLMS, aimed at effective zero-shot cross-lingual transfer. XeroAlign is an auxiliary training objective that is easily integrated into the unaltered primary task/model. XeroAlign leverages translated training data to bring the sentence embeddings in different languages closer together. We evaluated the XeroAligned XLM-R models (named XLM-RA) on zero-shot cross-lingual text classification, adversarial paraphrase detection and slot filling tasks, achieving SOTA (or near-SOTA) scores across four datasets covering 11 unique languages. Our ultimate vision is a level of zero-shot performance at or near that of Target models. The XeroAligned XLM-R partially achieved that goal by exceeding the intent classification and paraphrase detection accuracies of XLM-R trained with labelled data. We hope our work will inspire further progress in cross-lingual transfer learning.

Acknowledgements

We would like to thank the reviewers for their thoughtful comments and constructive feedback. We also want to thank the MindSpore⁷⁸ team for the technical support.

⁷<https://github.com/mindspore-ai>

⁸<https://mindspore.cn/>

References

- Suzanna Becker and Geoffrey E Hinton. 1992. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356):161–163.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2020. In-foxlm: An information-theoretic framework for cross-lingual language model pre-training. *arXiv preprint arXiv:2007.07834*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485.
- 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.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of ibm model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648.
- Yuwei Fang, Shuohang Wang, Zhe Gan, Siqi Sun, and Jingjing Liu. 2020. Filter: An enhanced fusion method for cross-lingual language understanding. *arXiv preprint arXiv:2009.05166*.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *arXiv preprint arXiv:2007.01852*.
- Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- Beliz Gunel, Jingfei Du, Alexis Conneau, and Ves Stoyanov. 2020. Supervised contrastive learning for pre-trained language model fine-tuning. *arXiv preprint arXiv:2011.01403*.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Karl Moritz Hermann and Phil Blunsom. 2014. Multilingual models for compositional distributed semantics. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 58–68.
- Junjie Hu, Melvin Johnson, Orhan Firat, Aditya Siddhant, and Graham Neubig. 2020a. Explicit alignment objectives for multilingual bidirectional encoders. *arXiv preprint arXiv:2010.07972*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020b. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- Alankar Jain, Bhargavi Paranjape, and Zachary C. Lipton. 2019. Entity projection via machine translation for cross-lingual NER. 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 1083–1092, Hong Kong, China. Association for Computational Linguistics.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the nlp world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86. Citeseer.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. Mlqa: Evaluating cross-lingual extractive question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330.
- Haoran Li, Abhinav Arora, Shuohui Chen, Anchit Gupta, Sonal Gupta, and Yashar Mehdad. 2021. **MTOP: A comprehensive multilingual task-oriented**

- semantic parsing benchmark. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2950–2962, Online. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. 2020. Xglue: A new benchmark dataset for cross-lingual pre-training, understanding and generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018.
- Zihan Liu, Genta Indra Winata, Zhaojiang Lin, Peng Xu, and Pascale Fung. 2020. Attention-informed mixed-language training for zero-shot cross-lingual task-oriented dialogue systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8433–8440.
- Fuli Luo, Wei Wang, Jiahao Liu, Yijia Liu, Bin Bi, Songfang Huang, Fei Huang, and Luo Si. 2020. Veco: Variable encoder-decoder pre-training for cross-lingual understanding and generation. *arXiv preprint arXiv:2010.16046*.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1659–1666.
- Lin Pan, Chung-Wei Hang, Haode Qi, Abhishek Shah, Mo Yu, and Saloni Potdar. 2020. Multilingual bert post-pretraining alignment. *arXiv preprint arXiv:2010.12547*.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. **Cross-lingual name tagging and linking for 282 languages**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. 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 H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Patti Price. 1990. Evaluation of spoken language systems: The atis domain. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Kunxun Qi and Jianfeng Du. 2020. Translation-based matching adversarial network for cross-lingual natural language inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8632–8639.
- Libo Qin, Minheng Ni, Yue Zhang, and Wanxiang Che. 2020. Cosda-ml: Multi-lingual code-switching data augmentation for zero-shot cross-lingual nlp. *arXiv preprint arXiv:2006.06402*.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. Cross-lingual transfer learning for multilingual task oriented dialog. 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 3795–3805.
- Leslie N Smith and Nicholay Topin. 2019. Superconvergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, page 1100612. International Society for Optics and Photonics.
- Shyam Upadhyay, Manaal Faruqui, Gokhan Tür, Hakkani-Tür Dilek, and Larry Heck. 2018. (almost) zero-shot cross-lingual spoken language understanding. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6034–6038. IEEE.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30:5998–6008.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019a. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in Neural Information Processing Systems*, 32.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In the Proceedings of ICLR.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, pages arXiv–1910.
- Weijia Xu, Batoool Haider, and Saab Mansour. 2020. End-to-end slot alignment and recognition for cross-lingual nlu. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5052–5063.

- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. Paws-x: A cross-lingual adversarial dataset for paraphrase identification. 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 3678–3683.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. Paws: Paraphrase adversaries from word scrambling. 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 1298–1308.
- Nanxuan Zhao, Zhirong Wu, Rynson WH Lau, and Stephen Lin. 2020. What makes instance discrimination good for transfer learning? *arXiv preprint arXiv:2006.06606*.
- Pierre Zweigenbaum, Serge Sharoff, and Reinhard Rapp. 2018. Overview of the third bucc shared task: Spotting parallel sentences in comparable corpora. In *Proceedings of 11th Workshop on Building and Using Comparable Corpora*, pages 39–42.