

BRIDGING THE DOMAIN GAP IN CROSS-LINGUAL DOCUMENT CLASSIFICATION

Anonymous authors

Paper under double-blind review

ABSTRACT

The scarcity of labeled training data often prohibits the internationalization of NLP models to multiple languages. Cross-lingual understanding has made progress in this area using language universal representations. However, most current approaches focus on the problem as one of aligning language and do not address the natural domain drift across languages and cultures. In this paper, We address the domain gap in the setting of semi-supervised cross-lingual document classification, where labeled data is available in a source language and only unlabeled data is available in the target language. We combine a state-of-the-art unsupervised learning method, masked language modeling pre-training, with a recent method for semi-supervised learning, Unsupervised Data Augmentation (UDA), to simultaneously close the language and the domain gap. We show that addressing the domain gap in cross-lingual tasks is crucial. We improve over strong baselines and achieve a new state-of-the-art for cross-lingual document classification.

1 INTRODUCTION

Recent advances in Natural Language Processing have enabled us to train high-accuracy systems for many language tasks. However, training an accurate system still requires a large amount of training data. It is inefficient to collect data for a new task and it is virtually impossible to annotate a separate data set for each language. To go beyond English and a few popular languages, we need methods that can learn from data in one language and apply it to others.

Cross-Lingual Understanding (XLU) has emerged as a field concerned with learning models on data in one language and applying it to others. Much of the work in XLU focuses on the zero-shot setting, which assumes that labeled data is available in one source language (usually English) and not in any of the target languages in which the model is evaluated. The labeled data can be used to train a high quality model in the source language. One then relies on general domain parallel corpora and monolingual corpora to learn to ‘transfer’ from the source language to the target language. Transfer methods can explicitly rely on machine translation models built from such parallel corpora. Alternatively, one can use such corpora to learn language universal representations to produce features to train a model in one language, which one can directly apply to other languages. Such representations can be in the form of cross-lingual word embeddings, contextual word embeddings, or sentence embeddings (Ruder et al. (2017); Lample & Conneau (2019); Schwenk & Douze (2017)). Using such techniques, recent work has demonstrated reasonable zero-shot performance for cross-lingual document classification (Schwenk & Li (2018)) and natural language inference (Conneau et al. (2018)).

What we have so far described is a simplified view of XLU, which focuses solely on the problem of aligning languages. This view assumes that, if we had access to a perfect translation system, and translated our source training data into the target language, the resulting model would perform as well as if we had collected a similarly sized labeled dataset directly in our target language. Existing work in XLU to date also works under this assumption. However, in real world applications, we must also bridge the domain gap across different languages, as well as the language gap. No task is ever identical in two languages, even if we group them under the same label, e.g. ‘news document classification’ or ‘product reviews sentiment analysis’. A Chinese customer might express sentiment differently than his American counterpart. Or French news might simply cover different topics than

English news. As a result, any approach which ignores this domain drift will fall short of native in-language performance in real world XLU.

In this paper, we propose to jointly tackle both language and domain transfer. We consider the semi-supervised XLU task, where in addition to labeled data in a source language, we have access to unlabeled data in the target language. Using this unlabeled data, we combine the aforementioned cross-lingual methods with recently proposed unsupervised domain adaptation and weak supervision techniques on the task of cross-lingual document classification (XLDC). In particular, we focus on two approaches for domain adaptation. The first method is based on masked language model (MLM) pre-training (as in Devlin et al. (2018)) using unlabeled target language corpora. Such methods have been shown to improve over general purpose pre-trained models such as BERT in the weakly supervised setting (Lee et al. (2019); Han & Eisenstein (2019)). The second method is unsupervised data augmentation (UDA) (Xie et al. (2019)), where synthetic paraphrases are generated from the unlabeled corpus, and the model is trained on a label consistency loss.

While both of these techniques were proposed previously, in both cases there are some open questions when applying them on the cross-lingual problems. For instance when performing data augmentation, one could generate augmented paraphrases in either the source or the target language or both. We experiment with various approaches and provide guidelines with ablation studies. Furthermore, we find that the value of additional labeled data in the source language is limited due to the train-test discrepancy of XLDC tasks. We propose to alleviate this issue by using self-training technique to do the domain adaptation from the source language into the target language. By combining these methods, we are able to reduce error rates by an average 44% over a strong XLM baseline, setting a new state-of-the-art for cross-lingual document classification.

2 RELATED WORK

2.1 CROSS-LINGUAL UNDERSTANDING

Cross-lingual document classification was first introduced in (Bel et al., 2003). The subsequent work (Prettenhofer & Stein, 2010) proposes the cross-lingual sentiment classification datasets, and (Lewis et al., 2004; Klementiev et al., 2012; Schwenk & Li, 2018) have extended this to the news domain. Cross-lingual understanding has also been applied to other NLP tasks, with datasets available in dependency parsing (Nivre et al., 2016), natural language inference (XNLI) (Conneau et al., 2018) and question answering (Liu et al. (2019)).

Cross-lingual methods gained popularity with the advent of cross-lingual word embeddings (Mikolov et al. (2013)). Since then, many methods have been proposed to better align the word embedding spaces of different languages (see Ruder et al. (2017) for a survey). Recently, more sophisticated extensions have been proposed based on seq2seq training of cross-lingual sentence embeddings (Schwenk & Douze (2017); Artetxe & Schwenk (2018)) and contextual word embeddings pre-trained on masked language modeling, notably multilingual BERT (Devlin et al. (2018)) and the cross-lingual language model (XLM) of Lample & Conneau (2019). We use XLM as our baseline representation in all experiments, as it’s the current state-of-the-art on the commonly used XNLI benchmark for cross-lingual understanding.

2.2 DOMAIN ADAPTATION

Domain adaptation, closely related to transfer learning, has a rich history in machine learning and natural language processing (Pan & Yang (2009)). Such methods have long been applied to document classification tasks (Blitzer et al., 2007; Glorot et al., 2011; Al-Moslmi et al., 2017; Xu & Yang, 2017).

Domain adaptation for NLP is intimately related to transfer learning and semi-supervised learning (Chapelle et al. (2009)). Transfer learning has made tremendous advances recently due to the success of pre-training representations using language modeling as a source task (Radford et al. (2018); Peters et al. (2018); Devlin et al. (2018)). While such representations trained on large amounts of general domain text have been shown to transfer well generally, performance still suffers when the target domain is sufficiently different than what the models were pre-trained on. In such cases, it is known that further pre-training the language model on in-domain text is helpful (Howard & Ruder

(2018); Chronopoulou et al. (2019)). It is natural to use unsupervised domain data for this task, when available (Lee et al. (2019); Han & Eisenstein (2019)).

The study of weakly supervised learning in language processing is relatively new (Johnson & Zhang, 2016; Yu et al., 2018). Most recently, Xie et al. (2019) has introduced an unsupervised data augmentation (UDA) technique to demonstrate improvements in the few-shot learning setting. Here, we extend this technique to facilitate cross-lingual and cross-domain transfer.

3 FRAMEWORK

In this section, we formally define the problem discussed in this paper and describe the proposed approach in detail.

3.1 PROBLEM FORMULATION

In vanilla zero-shot cross-lingual document classification, it is assumed that we have available a labeled dataset in a source language (English in our case), which we can denote by $L_{src} = \{(x, y) | x \sim P_{src}(x)\}$, where $P_{src}(x)$ is the prior distribution of task data in the source language. It is assumed that no data is available for the task in the target language. General purpose parallel and monolingual resources are used to train a cross-lingual classifier on the labeled source data, which is then applied to the target language data at test time.

In this work, we also assume access to a large unlabeled corpus in the target language, $U_{tgt} = \{x | x \sim P_{tgt}(x)\}$, which is usually the case in practical applications. We aim to utilize this domain-specific unlabeled data to help model be adapted to the target domain and gain the better generalization ability. We refer to this setting as *semi-supervised XLDC*, although we’re still in the zero-shot setting, in that no labeled data is used in the target language.

3.2 BASELINE APPROACHES

There are two standard ways to transfer knowledge across the languages in the vanilla zero-shot setting: (1) using a translation system to translate the labeled samples, such as translate-train and translate-test methods, and (2) learning a multilingual embedding system to obtain a language irrelevant representations of the data. In this paper, we adopt the second approach as the basic model, and utilize the XLM model (Lample & Conneau, 2019) as our base model, which has been pre-trained by large-scale parallel and monolingual data from various languages. Because XLM is a multilingual embedding system, a baseline is obtained by fine-tuning XLM with the labeled set $L_{src}(x)$ and directly applying the resulting model to the target language. In the experiments section, we also discuss the combination of the XLM and the translation based approaches.

3.3 SEMI-SUPERVISED XLDC

As argued in Introduction, even with a perfect translation or multilingual embedding system, we still face the domain-mismatch problem. This mismatch may limit the generalization ability of the model during testing time. To fully adapt the classifier to the target distribution, we explore the following approaches, each of which leverages unlabeled data in the target language in different ways.

Masked Language Model pre-training BERT (Devlin et al., 2018) and its derivations (such as XLM) are trained on general domain corpora. A standard practice is to further pre-train to adapt to a particular domain when data is available. This technique can let model learn the prior knowledge from the target domain and lead to better performance. We refer to this approach as the MLM pre-training.

However, in the cross-lingual setting, fine-tuning the XLM model in the target language can make the model degenerate in the source language, decreasing its ability to transfer across languages. This problem prohibits the fine-tuning baseline to learn a reasonable model. Therefore, in this case, we take care to use this method in combination with the translate-train method, which translates all labeled samples into the target language.




Figure 1: The process diagram of the UDA with self-training method

Unsupervised Data Augmentation The second approach is utilizing the state-of-the-art semi-supervised learning technique, Unsupervised Data Augmentation (UDA) algorithm (Xie et al., 2019), to leverage the unlabeled data. The objective function of UDA can be written as,

$$\min_{\theta} \mathbb{E}_{(x,y) \in L_{src}} [-\log(p_{\theta}(y|x))] + \lambda \mathbb{E}_{x \in U_{tgt}} [D_{KL}(p_{\theta}(y|\hat{x})||p_{\theta}(y|x))] \quad (1)$$

where \hat{x} is an augmented sample generated by a predefined augmentation function $\hat{x} = q(x)$. The augmentation function can be a paraphrase generation model, or a noising heuristic. Here, we use a machine translation system as the paraphrase generation model. The UDA loss enforces the classifier to produce label consistent predictions for pairs of original and augmented samples. With UDA method, the model is better adapted to the target domain by learning a label consistent model over the target domain.

In the cross-lingual setting, there are multiple ways of generating augmented samples using translation. One could translate samples from the target language into the source language and use this cross-lingual pair as the augmented sample. Alternatively, one could translate back into the target language and use only target-language augmented samples. We find that the latter works best. It is also possible to do data augmentation using source domain unlabeled data. The results of these comparisons are included in our detailed ablation study in the experiments section.

Alleviating the Train-Test Discrepancy of the UDA Method With the UDA algorithm, the classifier is able to explore the prior information on the target domain, however it still suffers from the train-test discrepancy. During the testing phase, our goal is to maximize the classifier performance on the real samples in target language, which focus on the samples from the distribution $P_{tgt}(x)$. Upon observing the training objective of the UDA method, Eq. (1), one can see that the data x that feed to model in the training phrase is sampled from three domains: (1) the source domain $P_{src}(x)$, (2) the target domain $P_{tgt}(x)$ and (3) the augmented sample domain $P_{aug}(x)$. On the other hand, the testing phrase only processes data from the target domain $P_{tgt}(x)$.

The source and target domain are mismatched, due to differences in language as argued earlier. Furthermore, the augmented domain, although generated from the target domain, can also be mismatched, due to artifacts introduced by the translation system. This can be especially problematic, since the UDA method needs diversity in the augmented samples to perform well (Xie et al., 2019), which trades off against their quality.

We propose to apply the self-training technique (Lee, 2013) to tackle this problem. We first train a classifier based on the UDA algorithm and denote it as $f^*(x)$, which is the teacher model used to score the unlabeled data U_{tgt} from the target domain. Then we fine-tune a new XLM model using the soft classification loss function with the pseudo-labeled data, which is written as,

$$\min_{\theta} \mathbb{E}_{x \in U_{tgt}} D_{KL}(f_{\theta}(x)||f^*(x)) \quad (2)$$

Following this process, we obtain a new classifier trained only based on the target domain, which does not suffer from the train-test mismatch problem. We show that this process provides better generalization ability compared to the teacher model. A process diagram of this method is presented in figure 1.

4 EXPERIMENTS

In this section, we present a comprehensive study on two benchmark tasks, cross-lingual sentiment classification, and cross-lingual news classification.

4.1 DATASETS

Sentiment Classification In this task, we test the proposed framework on a sentiment classification benchmark in three target languages, i.e. French, German and Chinese. The French, German and English data come from the benchmark cross-lingual Amazon reviews dataset (Prettenhofer & Stein, 2010), which we denote as *amazon-fr*, *amazon-de* and *amazon-en*. We merge training and testing samples from all product categories in one language, which leads to 6000 training samples. However, for the purpose of facilitating fair comparison with previous work, we also provide results for category-wise settings to compare previous state-of-the-art performance. The number of unlabeled samples from *amazon-fr* is 54K, and from *amazon-de* is 310K.

For Chinese, we use the Chinese Amazon (*amazon-cn*) (Zhang et al., 2015b) and Dianping (Zhang et al., 2014) datasets. Dianping is a business review website similar to Yelp. The training data for *amazon-cn* is *amazon-en*, and for *dianping* it is the Yelp dataset (Zhang et al., 2015a). In these two cases, the size of the training sample is 2000. For both *amazon-cn* and *dianping* datasets, we have 4M unlabeled examples. Because the number of the unlabeled set is very large, we randomly sample 10% for the UDA algorithm.

News Classification We use the MLDoc dataset (Schwenk & Li, 2018) for this task. The MLDoc dataset is a subset of RCV2 multilingual news dataset (Lewis et al., 2004). It has 4 categories, i.e. Corporate/Industrial, Economics, Government/Social and Markets, and each category has 250 training samples. We use the rest of the news documents in RCV2 dataset as the unlabeled data. The number of unlabeled samples for each language ranges from 5K to 20K, which is relatively smaller compared to the sentiment classification task.

Because the XLM model is pre-trained on 15 languages, we ignore languages which are not supported by XLM in the above benchmark datasets.

The pre-processing scripts for the above datasets, augmented samples, pretrained models and experiment settings needed to reproduce results will be released in the Github repo.

4.2 MASKED LANGUAGE MODEL PRE-TRAINING STRATEGY

As introduced in section 3.3, we apply MLM pre-training on the unlabeled data corpus to obtain a domain-specific XLM, denoted as XLM_{ft} in the following sections. The pre-training strategies for the two tasks are slightly different. In the sentiment classification task, because the size of the unlabeled corpus in each target domain is large enough, we fine-tune an XLM with MLM loss for each target domain respectively. In contrast, we do not have enough unlabeled data in each language in the MLDoc dataset, therefore we integrate unlabeled data from all languages as the training corpus. As a result the XLM_{ft} still preserves its language universality in this task.

4.3 MAIN RESULTS

We compare the following models:

- Fine-tune (Ft): Fine-tuning the pre-trained model with the source-domain training set. In the case of XLM_{ft} , the training set is translated into the target language.
- Fine-tune with UDA (UDA): This method utilizes the unlabeled data from the target domain by optimizing the UDA loss function (Eq. (1)).
- Self-training based on the UDA model (UDA+Self): We first train the Ft model and UDA model, and choose the one with better develop accuracy as the teacher model. The teacher model is used to train a new XLM student using only unlabeled data U_{tgt} in the target domain, as described above.

We report the results of applying these three methods on both the original XLM model and the XLM_{ft} model. In order to keep the notation simple, we use parenthesis after the method name to

indicate which basic model was used, such as UDA(XLM_{ft}). The details about the implementation and hyper-parameter tuning are included in Appendix A.1.

The results for the cross-lingual sentiment classification task are summarized in table 1. As our experiment setting on the cross-lingual amazon dataset is different from previous publications, in order to provide a fair comparison with previous works, we summarize the results of the standard category-wise setting in table 3. The results for cross-lingual news classification is included in table 2. The last column “Unlabeled” in these tables indicates whether this method utilizes the unlabeled data. For the monolingual baselines, the models are trained with labeled data from the target domain. The size of the labeled set is the same as the English training set used for cross-lingual experiments.

We can summarize our findings as follows:

- Looking at Ft(XLM) results, it is clear that without the help of unlabeled data from the target domain, there still exists a substantial gap between the model performance of the cross-lingual settings and the monolingual baselines, even when using state-of-the-art pre-trained cross-lingual representations.
- Both the UDA algorithm and MLM pre-training can offer significant improvements by utilizing the unlabeled data. In the sentiment classification task, where the unlabeled data size is larger, Ft(XLM_{ft}) model using MLM pre-training consistently provides larger improvements compared with the UDA method. On the other hand, the MLM method is relatively more resource intensive and takes longer to converge (see Appendix A.5). In contrast, in the MLdoc dataset, when the size of the unlabeled samples is limited, the UDA method is more helpful.
- The combination of both methods - as in the UDA(XLM_{ft}) model - consistently outperforms either method alone. In this case the additional improvement provided by the UDA algorithm is smaller, but still consistent.
- In the sentiment classification task, we observe the self-training technique consistently improves over its teacher model. It offers best results in both XLM and XLM_{ft} based classifiers. The results demonstrate that self-training process is able to alleviate the train-test distribution mismatch problem and provide better generalization ability. In the MLdoc dataset, self-training also achieves the best results overall, however the gains are less clear. We hypothesize that this technique is not as useful without enough number of unlabeled samples.
- Finally, comparing with the best cross-lingual results and monolingual fine-tune baseline, we are able to completely close the performance gap by utilizing unlabeled data in the target language. Furthermore, our framework reaches new state-of-the-art results, improving over vanilla XLM baselines by 44% on average.

Furthermore, we provide an additional baseline, which only uses English samples to perform semi-supervised learning, whose details are in Appendix A.2. The experiment results show that it lags behind the ones using unlabeled data from the target domain. This observation also justifies the importance of information from the target domain in the XLDC task.

Base Model	Train Test	amazon-en amazon-fr	amazon-en amazon-de	amazon-en amazon-cn	yelp dianping	Unlabeled
XLM	Ft	11.35	12.75	15.5	12.66	✗
	UDA	9.83	10.08	11.92	8.37	✓
	UDA+Self	9.03	9.32	11.19	7.80	✓
XLM_{ft}	Ft	7.67	6.70	11.32	6.15	✓
	UDA	7.15	6.30	11.92	5.21	✓
	UDA+Self	6.67	5.77	9.46	4.80	✓
Monolingual Baselines						
XLM	Ft	8.32	10.67	11.69	8.73	✗
XLM_{ft}	UDA	5.95	6.12	7.74	4.64	✓

Table 1: Error rates for the sentiment classification task.

Base Model	Train Test	fr	de	en es	zh	ru	Unlabeled
	pre-XLM sota	21.97	13.75	20.7	25.27	33.22	✗
XLM	Ft	7.95	6.03	12.08	11.95	26.95	✗
	UDA	4.97	4.30	7.67	9.85	27.22	✓
	UDA+Self	4.75	4.35	7.78	9.70	27.68	✓
XLM _{ft}	Ft	6.30	6.43	8.70	11.02	25.78	✓
	UDA	4.65	4.63	6.97	7.15	16.50	✓
	UDA+Self	4.60	4.27	6.53	7.15	15.65	✓
Monolingual Baselines							
XLM	Ft	5.35	3.77	3.95	7.85	11.05	✗
XLM _{ft}	UDA	3.95	3.05	3.20	6.68	10.30	✓

Table 2: Error rates for news document classification. The pre-XLM sota results are provided by Artetxe & Schwenk (2018)

Base Model	dataset category	amazon-fr			amazon-de			Unlabeled
		book	dvd	music	book	dvd	music	
	pre-XLM sota	14.0	14.3	14.0	13.6	13.9	12.2	✓
XLM	UDA+Self	7.75	8.70	9.35	8.30	10.00	9.96	✓
XLM _{ft}	UDA+Self	6.50	7.05	7.15	4.80	7.10	5.55	✓

Table 3: Error rates for the sentiment classification task by product category. The pre-XLM sota results are provided by Chen et al. (2019).

4.4 LABELED DATA IN THE SOURCE LANGUAGE HAS LIMITED VALUE

In this section, we provide evidence for the train-test domain discrepancy in the context of the UDA method, by showing that adding more labeled data in the source language does not improve target task accuracy after a certain point. Figure 2 plots the target model performance vs. the number of labeled training samples in the cross-lingual and monolingual settings respectively. The figures are based on the UDA(XLM) method with 6 runs in the Yelp-Dianping cross-lingual setting. The dot is the average accuracy and the filling area contains one standard derivation.

We observe that, in the cross-lingual setting, the model performance peaks at around 10k training samples per category, and becomes worse with the larger training set. In contrast, the performance of the model improves consistently with more labeled data in the monolingual setting. This suggests that more training data from the source domain could harm model generalization ability in the target domain with UDA approach in the cross-lingual setting. In order to alleviate this issue, we propose to utilize the self-training technique, which abandons the data from the source domain and the augmentation domain, to maximize its performance in the target domain.

4.5 ABLATION STUDY: AUGMENTATION STRATEGIES

Next, we explore different augmentation strategies and their influence on the final performance. As stated in section 3.3, the augmentation strategy used in the main experiment is that we first translate the samples into English and translate them back to its original language. We refer to this strategy as augmenting "from target domain to target domain" and abbreviate it as **t2t**.

We also explore two additional augmentation strategies: (1) First, we do not translate the samples back to the target language and directly use English samples as the augmented samples, denoted as **t2s**. Naturally, the parallel samples in two languages have the same sentiment information and different input format which are suitable to be used as the augmentation sample pairs for the multilingual system such as XLM. (2) The second approach is to leverage unlabeled data from other language




Figure 2: Plot of the accuracy vs. the number of training samples in log scale when running UDA(XLM) method. On the left is the cross-lingual setting. On the right is the monolingual setting.

domains. Here, we attempt to use the English unlabeled data. We translate them into the target language as the augmented samples. This strategy is denoted as **s2t**.

train	amazon-fr	amazon-en	amazon-cn	MLDoc en				
test		amazon-de		fr	de	es	zh	ru
t2t	9.83	10.08	11.92	4.65	4.63	6.97	8.15	16.5
t2s	10.97	13.15	15.07	5.03	5.00	9.38	8.35	27.90
t2t + t2s	9.82	12.07	11.93	5.07	4.33	7.15	8.90	28.60
s2t	9.32	10.38	12.28	5.53	6.25	7.73	8.50	29.17
t2t + s2t	8.98	10.87	11.79	4.85	4.40	7.75	8.27	30.33

Table 4: Error rates when using different augmentation strategies and their combinations. Results for sentiment classification shown on the left, and news document classification on the right.

Table 4 summarizes the performance of the proposed augmentation strategies and their combinations with the UDA(XLM) method in the sentiment classification and the UDA(XLM_{ft}) in the news classification settings. From the results, we conclude that **t2t** is the best performing approach, as it’s the best matched to the target domain. Leveraging the unlabeled data from other domains does not offer consistent improvement, however can provide additional value in isolated cases.

We include additional ablations regarding translation system in the appendix, including the application of translate-train method in our experiments (section A.3) and effects of hyper-parameters (section A.4), and finally we discuss the application of our framework in the monolingual cross-domain document classification problem setting in Appendix B.

5 CONCLUSION

In this paper, we tackled the domain mismatch challenge in cross-lingual document classification - an important, yet often overlooked problem in cross-lingual understanding. We provided evidence for the existence and importance of this problem, even when utilizing strong pre-trained cross-lingual representations.

We proposed a framework combining cross-lingual transfer techniques with three domain adaptation methods; unsupervised data augmentation, masked language model pre-training and self-training, which can leverage unlabeled data in the target language to moderate the domain gap. Our results show that by removing the domain discrepancy, we can close the performance gap between cross-lingual transfer and monolingual baselines almost completely for the document classification task. We are also able to improve the state-of-the-art in this area by a large margin. While document classification is by no means the most challenging task for XLU, we believe the strong gains that we demonstrated could be extended to other cross-lingual tasks, such as cross-lingual question answering and event detection. Developing cross-lingual methods which are competitive with in-language models for real world, semantically challenging NLP problems remains an open problem and subject of future research.

REFERENCES

- Tareq Al-Moslmi, Nazlia Omar, Salwani Abdullah, and Mohammed Albared. Approaches to cross-domain sentiment analysis: A systematic literature review. *IEEE Access*, 5:16173–16192, 2017.
- Ahmed Aly, Kushal Lakhotia, Shicong Zhao, Mrinal Mohit, Barlas Oguz, Abhinav Arora, Sonal Gupta, Christopher Dewan, Stef Nelson-Lindall, and Rushin Shah. Pytext: A seamless path from nlp research to production. *arXiv preprint arXiv:1812.08729*, 2018.
- Mikel Artetxe and Holger Schwenk. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *arXiv preprint arXiv:1812.10464*, 2018.
- Nuria Bel, Cornelis HA Koster, and Marta Villegas. Cross-lingual text categorization. In *International Conference on Theory and Practice of Digital Libraries*, pp. 126–139. Springer, 2003.
- John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, pp. 440–447, 2007.
- Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o. et al., eds., 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009.
- Zhenpeng Chen, Sheng Shen, Ziniu Hu, Xuan Lu, Qiaozhu Mei, and Xuanzhe Liu. Emoji-powered representation learning for cross-lingual sentiment classification. In *The World Wide Web Conference*, pp. 251–262. ACM, 2019.
- Alexandra Chronopoulou, Christos Baziotis, and Alexandros Potamianos. An embarrassingly simple approach for transfer learning from pretrained language models. *arXiv preprint arXiv:1902.10547*, 2019.
- Alexis Conneau, Guillaume Lample, Ruty Rinott, Adina Williams, Samuel R Bowman, Holger Schwenk, and Veselin Stoyanov. Xnli: Evaluating cross-lingual sentence representations. *arXiv preprint arXiv:1809.05053*, 2018.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 513–520, 2011.
- Xiaochuang Han and Jacob Eisenstein. Unsupervised domain adaptation of contextualized embeddings: A case study in early modern english. *arXiv preprint arXiv:1904.02817*, 2019.
- Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*, 2018.
- Rie Johnson and Tong Zhang. Supervised and semi-supervised text categorization using lstm for region embeddings. *arXiv preprint arXiv:1602.02373*, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. Inducing crosslingual distributed representations of words. In *Proceedings of COLING 2012*, pp. 1459–1474, 2012.
- Guillaume Lample and Alexis Conneau. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*, 2019.
- Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*, volume 3, pp. 2, 2013.

- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: pre-trained biomedical language representation model for biomedical text mining. *arXiv preprint arXiv:1901.08746*, 2019.
- David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397, 2004.
- Jiahua Liu, Yankai Lin, Zhiyuan Liu, and Maosong Sun. Xqa: A cross-lingual open-domain question answering dataset. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pp. 2358–2368, 2019.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*, 2013.
- 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. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pp. 1659–1666, 2016.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 2018.
- Peter Prettenhofer and Benno Stein. Cross-language text classification using structural correspondence learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pp. 1118–1127, 2010.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf, 2018.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. A survey of cross-lingual word embedding models. *arXiv preprint arXiv:1706.04902*, 2017.
- Holger Schwenk and Matthijs Douze. Learning joint multilingual sentence representations with neural machine translation. *arXiv preprint arXiv:1704.04154*, 2017.
- Holger Schwenk and Xian Li. A corpus for multilingual document classification in eight languages. *arXiv preprint arXiv:1805.09821*, 2018.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation. *arXiv preprint arXiv:1904.12848*, 2019.
- Ruochen Xu and Yiming Yang. Cross-lingual distillation for text classification. *arXiv preprint arXiv:1705.02073*, 2017.
- Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauero, Haoyu Wang, and Bowen Zhou. Diverse few-shot text classification with multiple metrics. *arXiv preprint arXiv:1805.07513*, 2018.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pp. 649–657, 2015a.
- Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pp. 83–92. ACM, 2014.

Yongfeng Zhang, Min Zhang, Yi Zhang, Guokun Lai, Yiqun Liu, Honghui Zhang, and Shaoping Ma. Daily-aware personalized recommendation based on feature-level time series analysis. In *Proceedings of the 24th international conference on world wide web*, pp. 1373–1383. International World Wide Web Conferences Steering Committee, 2015b.

A APPENDIX OF CROSS-LIGUAL EXPERIMENTS

A.1 IMPLEMENTATION DETAILS

The experiments in this paper are based on the PyTorch (Paszke et al., 2017) and Pytext (Aly et al., 2018) package. We use the Adam (Kingma & Ba, 2014) as the optimizer. For all experiments, we grid search the learning rate in the set $\{5 \times 10^{-6}, 1 \times 10^{-5}, 2 \times 10^{-5}\}$. When using UDA method, we also try the three different annealing strategies introduced in the UDA paper (Xie et al., 2019), and the λ in (1) is always set as 1. The batch size in the Ft and UDA+Self method is 128. In the UDA method, the batch size is 16 for the labeled data and 80 for the unlabeled data. Due to the limitation of the GPU memory, in all experiments, we set the length of samples as 256, and cut the input tokens exceeding this length. Finally, we report the results with the best hyper-parameters.

As for the augmentation process, we sweep the temperature which controls the diversity of beam search in translation. The best temperature for “en-de, en-fr, en-es” and “en-ru” are 1.0 and 0.6, the sampling space is the whole vocabulary. In the “en-zh” setting, the temperature is 1.0 and the sampling space is the top 100 tokens in the vocabulary. We note that this uses the Facebook production translation models, and results could vary when other translation systems are applied. For reproducibility, we will release the augmented datasets that we generated.

A.2 ABLATION STUDY: THE BASELINE WITH ENGLISH UNLABELED DATA

Here, we provide a baseline only using English samples to perform semi-supervised learning. More specifically, we first train the model with English unlabeled data and augmented samples, then tests it on different target domains. This approach is similar to the traditional translate-test method. This method offers a baseline, which merely increases the size of data but without providing the target domain information. During the test phrasing, we experiment with two input strategies. One is using the original test samples, and another is translating the samples into English. We report the results (Table 5) of the UDA(XLM) method with two input strategies and compare them with the main results, which uses the unlabeled data from the target domain. First, we observe that the performance of using original and translated samples is similar. Second, compared with Ft(XLM) baselines in section 4.3, utilizing the unlabeled data from the English domain is slightly better than only training with labeled data, but it still lags behind the performance of using the unlabeled data from the target domain.

lang of test	Unlabeled	amazon-fr	amazon-de	amazon-cn		fr	de	es	zh	ru
original	English	10.57	12.98	13.73		6.47	6.13	12.35	11.6	27.52
translated	English	10.68	12.40	15.21		6.20	6.20	11.88	10.12	31.37
original	Target	9.83	10.08	11.92		4.97	4.30	8.67	9.85	27.22

Table 5: The first part is the baseline results using the English unlabeled data. The second part is the results using the unlabeled data from the target domain, which are copied from the section 4.3.

A.3 ABLATION STUDY: TRANSLATE-TRAIN

As discussed earlier, fine-tuning XLM on the target language would depreciate the multilingual ability of the model. We apply the translate-train method to tackle this problem. In order to understand the influence of this strategy when using the proposed framework, we perform an ablation study. We test 3 input strategies: (1) English: use the original English data as training data. (2) tr-train: use the translate-train strategy, which translate the training data into the target language. (3) both: we combine the (1) and (2) as the training data. We report the results of the UDA(XLM) method in the sentiment classification tasks and UDA(XLM_{ft}) method in the news classification tasks in table 6.

train test	amazon-en			MLDoc en				
	amazon-fr	amazon-de	amazon-cn	fr	de	es	zh	ru
English	9.83	10.08	12.02	4.75	4.63	8.07	8.15	28.28
tr-train	9.92	11.70	15.31	5.57	5.00	6.97	9.20	16.50
both	9.75	11.02	16.02	4.95	4.87	6.33	8.35	28.63

Table 6: Ablation Study about the translate-train strategies. Results for sentiment classification shown on the left, and news document classification on the right.

We observe that in most cases, using the original training examples achieves the best performance. However, in special cases such as MLDoc-ru, the translate-train method achieves better performance. We recommend trying both approaches in practice.

A.4 ABLATION STUDY: DIVERSITY IN BEAM DECODING OF DATA AUGMENTATION

Given a translation system, we use the sample decoding strategy to translate the sample. The sample space is the entire vocabulary space. We tune the temperature of μ of the softmax distribution. As discussed in Xie et al. (2019), this controls the trade-off between quality and diversity. When $\mu = 0$, the sampling reduces to the greedy search and produce the best quality samples. When $\mu = 1$, the sampling produces diverse outputs but loses some semantic information. The table A.4 illustrates the influence of μ value to the final performance in the English-to-French and English-to-German settings. The results show that temperature has a significant influence on the final performance. However, because the quality of translation systems for different language pairs are not the same, their best temperature also varies. In the appendix A.1, we include the best temperature values for the translation systems used in this paper.

μ	amazon-fr	amazon-de	mldoc-fr	mldoc-de
0.6	11.92	13.27	8.95	7.40
0.8	9.63	10.6	5.55	4.70
1.0	9.83	10.08	4.97	4.30
1.2	10.83	13.58	7.87	5.50

Table 7: Effect of the temperature of the translation sampling decoder.

A.5 COMPUTATION TIME OF UDA AND MLM PRETRAINING

From the main results in section 4.3, we can see that MLM pre-training can offer better improvements over the UDA method. However, it is also more resource intensive, since MLM pre-training is a token level task with a large output space, which leads to more computationally intensive updates and also takes longer to converge. In our experiments, we used NVIDIA V100-16G GPUs to train all models. 8 GPUs were used to train Ft and UDA methods, and 32 GPUs to perform MLM pretraining. In the "amazonen->amazonfr" setting, for example, the unlabeled set contains 50K unlabeled samples and 8M tokens after BPE tokenization. The Ft method takes 3.2 GPU hours to converge. The UDA method training takes 16.8 GPU hours, excluding the time it takes to generate augmented samples (which we handle as part of data pre-processing). MLM pre-training takes upwards of 500 GPU hours to converge. This is another factor which should be taken into account.

B MONOLINGUAL DOMAIN ADAPTATION

As further evidence that our method addresses the domain mismatch, we apply our framework to the monolingual cross-domain document classification problem. We again focus on sentiment classification where data comes from two different domains, product reviews (amazon-en, amazon-cn) and business reviews (Yelp and Dianping). We train and test on the same language, only transferring across domains. We consider the two domain-pairs, amazonen-yelp and amazoncn-dianping. The results are illustrated in table 8. Conclusions are similar to the cross-domain setting (section 4.3):

- There exists a clear gap between the cross-domain and in-domain results of the Ft method, even when using strong pre-trained representations.
- By leveraging the unlabeled data from the target domain, we can significantly boost the model performance.
- Best results are achieved with our combined approach and almost completely matches the in-domain baselines.

Base Model	Train Test	amazon-en yelp	yelp amazon-en	amazon-cn dianping	dianping amazon-cn	Unlabeled
XLM	Ft	8.46	14.52	12.86	14.96	✗
	UDA	5.43	10.10	8.31	11.19	✓
	UDA+Self	4.94	9.58	8.71	10.69	✓
XLM _{ft}	Ft	4.51	10.80	7.06	9.96	✓
	UDA	3.78	7.87	5.54	7.38	✓
	UDA+Self	3.36	7.47	5.34	7.40	✓
In-Domain Baselines						
XLM	Ft	5.80	10.90	11.43	12.31	✗
XLM _{ft}	UDA	3.34	7.57	4.64	7.74	✓

Table 8: Error rates for cross-domain document classification