

The RELX Dataset and Matching the Multilingual Blanks for Cross-Lingual Relation Classification

Abdullatif Köksal

Department of Computer Engineering
Boğaziçi University

abdullatif.koksal@boun.edu.tr

Arzucan Özgür

Department of Computer Engineering
Boğaziçi University

arzucan.ozgur@boun.edu.tr

Abstract

Relation classification is one of the key topics in information extraction, which can be used to construct knowledge bases or to provide useful information for question answering. Current approaches for relation classification are mainly focused on the English language and require lots of training data with human annotations. Creating and annotating a large amount of training data for low-resource languages is impractical and expensive. To overcome this issue, we propose two cross-lingual relation classification models: a baseline model based on Multilingual BERT and a new multilingual pretraining setup, which significantly improves the baseline with distant supervision. For evaluation, we introduce a new public benchmark dataset for cross-lingual relation classification in English, French, German, Spanish, and Turkish, called **RELX**. We also provide the **RELX-Distant** dataset, which includes hundreds of thousands of sentences with relations from Wikipedia and Wikidata collected by distant supervision for these languages. Our code and data are available at: <https://github.com/boun-tabii/RELX>

1 Introduction

Extracting useful information from unstructured text is one of the most essential topics in Natural Language Processing (NLP). Relation classification can help achieving this objective by enabling the automatic construction of knowledge bases and by providing useful information for question answering models (Xu et al., 2016). Given an entity pair (e_1, e_2) and a sentence S that contains these entities, the goal of relation classification is to predict the relation $r \in R$ between e_1 and e_2 from a set of predefined relations, which may include ‘no relation’ as well. For example, with the help of relation classification, we can create semantic triples such as (*Rocky Mountain High School*, *founded*,

1973) from a sentence like “*Rocky Mountain High School opened at its current location in 1973 and was expanded in 1994.*”, where ‘Rocky Mountain High School’ and ‘1973’ are the given entities and ‘founded’ is the relation between them based on this sample sentence.

Traditionally, relation classification methods rely on hand-crafted features (Kambhatla, 2004). Lately, pretrained word embeddings (Mikolov et al., 2013) with RNN-LSTM architecture (Zhang and Wang, 2015; Xu et al., 2015) or transformers based models (Soares et al., 2019) have gained more attention in this domain. Although non-English content on the web is estimated as over 40% (Upadhyay, 2019) and the number of multilingual text-corpora is increasing (Indurkha, 2015), recent studies on relation classification have generally focused on the English language. These supervised approaches for relation classification are not easily adaptable to other languages, since they require large annotated training datasets, which are both costly and time-consuming to create.

The challenge of creating manually labeled training datasets for different languages can be alleviated through cross-lingual NLP approaches. In cross-lingual relation classification, the objective is to predict the relations in a sentence in a target language, while the model is trained with a dataset in a source language, which may be different from the target language. For example, a cross-lingual relation classification model should be able to extract semantic triples such as (*CD Laredo*, *founded*, *1927*) from a Spanish sentence like “*CD Laredo fue fundado en 1927 con el nombre de Sociedad Deportiva Charlestown.*¹” for the given entities ‘CD Laredo’ and ‘1927’, even when the annotated training data is solely in English.

Thanks to multilingual pretrained transformer

¹English Translation: CD Laredo was founded in 1927 with the name “Sociedad Deportiva Charlestown”.

models like Multilingual BERT (mBERT) (Devlin et al., 2019) and XLM (Conneau and Lample, 2019), cross-lingual models have been studied in depth for several NLP tasks such as question answering (Artetxe et al., 2020; Liu et al., 2019; Conneau et al., 2020), natural language inference (Conneau and Lample, 2019; Conneau et al., 2020; Wu and Dredze, 2019), and named entity recognition (Conneau et al., 2020).

In this paper, we first propose a baseline cross-lingual model for relation classification based on the pretrained mBERT model². Then, we introduce an approach called **Matching the Multilingual Blanks** to improve the relation classification ability of mBERT in different languages with the help of a considerable number of relation pairs collected by distant supervision. Prior works on cross-lingual relation classification use additional resources in the target language such as aligned corpora (Kim and Lee, 2012), machine translation systems (Faruqui and Kumar, 2015), or bilingual dictionaries (Ni and Florian, 2019). Our mBERT baseline model does not require any additional resources in the target language. The Matching the Multilingual Blanks model improves mBERT by utilizing the already available Wikipedia and Wikidata resources with distant supervision.

We present two new datasets for cross-lingual relation classification, namely RELX and RELX-Distant. **RELX** has been developed by selecting a subset of the commonly-used KBP-37 English relation classification dataset (Zhang and Wang, 2015) and generating human translations and annotations in the French, German, Spanish, and Turkish languages. The resulting dataset contains 502 parallel test sentences in five different languages with 37 relation classes. To our knowledge, RELX is the first parallel relation classification dataset, which we believe will serve as a useful benchmark for evaluating cross-lingual relation classification methods. **RELX-Distant** is a multilingual relation classification dataset collected from Wikipedia and Wikidata through distant supervision for the aforementioned five languages. We gather from 50 thousand upto 800 thousand sentences, whose entities have been labeled by the editors of Wikipedia. The relations among these entities are extracted from Wikidata.

Our contributions can be summarized as follows:

1. We introduce the **RELX** dataset, a novel

²<https://github.com/google-research/bert>

cross-lingual relation classification benchmark with 502 parallel sentences in English, French, German, Spanish, and Turkish.

2. To support distantly supervised models, we introduce the **RELX-Distant** dataset, which has hundreds of thousands of sentences with relations collected from Wikipedia and Wikidata for the mentioned five languages.
3. We first present a baseline mBERT model for cross-lingual relation classification and then, propose a novel multilingual distant supervision approach to improve the model.

The rest of the paper is organized as follows. The related work is discussed in Section 2. The details about the datasets are presented in Section 3. Our mBERT baseline model and the Matching the Multilingual Blanks (MTMB) procedure are described in Section 4. The experimental results for the mBERT model and MTMB are presented in Section 5. Finally, we draw conclusions and discuss future work in Section 6.

2 Related Work

In monolingual relation classification, traditional methods generally depend on hand-crafted features (Kambhatla, 2004). After the introduction of word embeddings (Mikolov et al., 2013; Pennington et al., 2014), many relation classification models used pretrained word embeddings with the RNN (Zhang and Wang, 2015; Xu et al., 2015) or CNN (Zeng et al., 2014; Nguyen and Grishman, 2015) architectures. With the strong performance of transformer networks for various NLP tasks (Devlin et al., 2019; Conneau and Lample, 2019; Peters et al., 2018), Soares et al. (2019) applied BERT with different representations to the relation classification task and showed the strength of it on several English datasets.

Before the introduction of multilingual transformers (Devlin et al., 2019; Conneau and Lample, 2019; Conneau et al., 2020), cross-lingual word embeddings have been widely used in zero-shot cross-lingual transfer with word embedding alignments for different tasks such as named entity recognition (Xie et al., 2018) and natural language inference (Conneau et al., 2018). This approach has also been utilized for cross-lingual relation classification (Ni and Florian, 2019). However, recently, multilingual deep transformers have attracted lots

of attention in several cross-lingual tasks such as question answering (Artetxe et al., 2020; Liu et al., 2019; Conneau et al., 2020), natural language inference (Conneau and Lample, 2019; Conneau et al., 2020; Wu and Dredze, 2019), and named entity recognition (Conneau et al., 2020). To the best of our knowledge, we present the first transformer based approach for the task of cross-lingual relation classification. In addition, we introduce a multilingual distant supervision method to improve the baseline transformer model. Soares et al. (2019) use a similar approach for monolingual relation classification, called Matching the Blanks. For the pretraining process, they collect pairs of English sentences based on the shared entities, annotated by an entity linking system. On the other hand, we propose a multilingual approach that utilizes Wikipedia and Wikidata, which are already available for many languages and have been successfully used for tasks such as multilingual question answering (Abdou et al., 2019) and named entity recognition (Nothman et al., 2013).

Most cross-lingual relation classification studies rely on parallel corpora, machine translation systems, or bilingual dictionaries. In (Kim et al., 2010; Kim and Lee, 2012), English labeled data are projected to Korean with parallel corpora to train relation classification models in Korean. Faruqui and Kumar (2015) apply a machine translation system to translate the sentence in a target language to a source language, so that a relation classification model trained with the source language can be used. Zou et al. (2018) make use of a Generative Adversarial Network (GAN) to transfer the feature representations from the source language to the target language with the help of machine translation systems. Ni and Florian (2019) employ bilingual word embedding mappings trained with bilingual dictionaries to develop a cross-lingual relation classification model.

In many studies, the multilingual ACE05 (Walker et al., 2006) relation classification dataset has been treated as cross-lingual for evaluation. ACE05 includes data in English, Arabic, and Chinese; however, it is not freely available, and the number of relations is rather small, which is 6. In (Ni and Florian, 2019), a relation classification dataset for 6 languages with 53 relation types has been used, yet this dataset is not publicly available. In this paper, we release the RELX dataset created with human annotations and the RELX-

Dataset	Total Sentences	Average Chars	Average Words
KBP-37			
Train	15917	181.21	30.28
Dev	1724	181.77	30.55
Test	3405	180.20	30.23
RELX			
English	502	171.18	28.88
French	502	186.63	30.99
German	502	188.27	27.73
Spanish	502	188.37	31.85
Turkish	502	170.76	23.60

Table 1: Comparative statistics of KBP-37 and RELX in different languages. Turkish translations have a lower number of words on average in the sentences due to the agglutinative nature of Turkish. The characters and words represent the average length of sentences in the corresponding dataset.

Distant dataset compiled using distant supervision. Both datasets are made publicly available for future cross-lingual relation classification studies.

3 The RELX and RELX-Distant Datasets

In this work, the training set of KBP-37 (Zhang and Wang, 2015) is used as a source in the English language for training. For evaluation, we introduce and make publicly available the RELX dataset in English, French, German, Spanish, and Turkish. We also present RELX-Distant, which we use for the pretraining procedure in the developed MTMB (Matching the Multilingual Blanks) approach, explained in Section 4.2.

3.1 RELX

We use the commonly-used KBP-37 English relation classification dataset for training due to its high amount of available training data. It contains 18 directional relations and a *no_relation* class, which results in 37 different classes. The statistics about KBP-37 are given in Table 1.

To create a cross-lingual relation classification benchmark, we selected a subset of 502 sentences from KBP-37’s test set by preserving the class distribution and the statistical features of KBP-37. 10,000 different subset selections are performed by conforming to the class distribution of KBP-37. The subset that is most similar to KBP-37 in terms of the sum of the normalized average character length and normalized average word length

English	<code><e1> Hoyte </e1> was born in <e2> Guyana </e2> ’s capital Georgetown.</code>
French	<code><e1> Hoyte </e1> est né à Georgetown, la capitale d’ <e2> Guyane </e2> .</code>
German	<code><e1> Hoyte </e1> wurde in der Hauptstadt Georgetown von <e2> Guyana </e2> geboren.</code>
Spanish	<code><e1> Hoyte </e1> nació en la capital de <e2> Guyana </e2> , Georgetown.</code>
Turkish	<code><e1> Hoyte </e1> , <e2> Guyana </e2> ’nın başkenti Georgetown’da doğdu.</code>
Category	<code>per:country_of_birth(e1,e2)</code>

Figure 1: Sample parallel sentences from RELX in different languages.

is chosen as the RELX dataset. Average character/word length normalization is performed by dividing to the average character/word length in the original KBP-37 test dataset. Due to the variety in the languages, the average number of characters and words in the sentences can differ for different languages, but the RELX-English and KBP-37 test set have similar distributions as summarized in Table 1. The average sentence length in the RELX-English dataset is slightly less than the KBP-37 test set, since we filtered problematic sentences that included URLs or consisted of more than one sentence.

The selected sentences are translated into French, German, Spanish, and Turkish by bilingual speakers who are advanced or native in both languages. They also marked the entities with (`<e1>`, `</e1>`) and (`<e2>`, `</e2>`) tags to match the same entities in these languages. Finally, professional translators from El Turco language services provider (eltur.co) performed language quality assessment for a randomly selected subset of RELX, containing 10% of the sentences from each language. Except article and synonym mistakes, there were less than three sentences with errors in each language and no critical errors were found in any of the translations. In Figure 1, we show an example of a parallel sentence from RELX with the marked entities for a sample relation.

3.2 RELX-Distant

We collected a large number of multilingual sentences with relations from Wikipedia and Wikidata by a distant supervision scheme (Mintz et al., 2009) and created the RELX-Distant weakly-labeled dataset for relation classification in English, French, German, Spanish, and Turkish.

The following steps are used to create RELX-Distant:

1. The Wikipedia dumps for the corresponding

Language	Number of Sentences
English	815689
French	652842
German	652062
Spanish	397875
Turkish	57114

Table 2: Number of sentences with a relation in each language in RELX-Distant.

languages are downloaded and converted into raw documents with Wikipedia hyperlinks in entities.

2. The raw documents are split into sentences with spaCy (Honnibal and Montani, 2017), and all hyperlinks, which refer to entities, are converted to their corresponding Wikidata IDs.
3. Sentences that include entity pairs with Wikidata relations (Vrandečić and Kröttsch, 2014) are collected.

The statistics about the created **RELX-Distant** dataset are provided in Table 2. After merging similar relations such as *capital* and *capital of*, RELX-Distant contains the following 24 relations, each of which include at least 1000 sentences in English Wikipedia.

author, capital, characters, continent, country of citizenship, country of origin, developer, ethnic group, father, instance of, language, located in country, member of, mother, owned by, parent organization, parent taxon, part of, partner, performer, place of birth, religion, sibling, spouse

4 Methods

Task Definition: In the cross-lingual relation classification task, we are given a source language dataset D_s with n_s sentences containing related

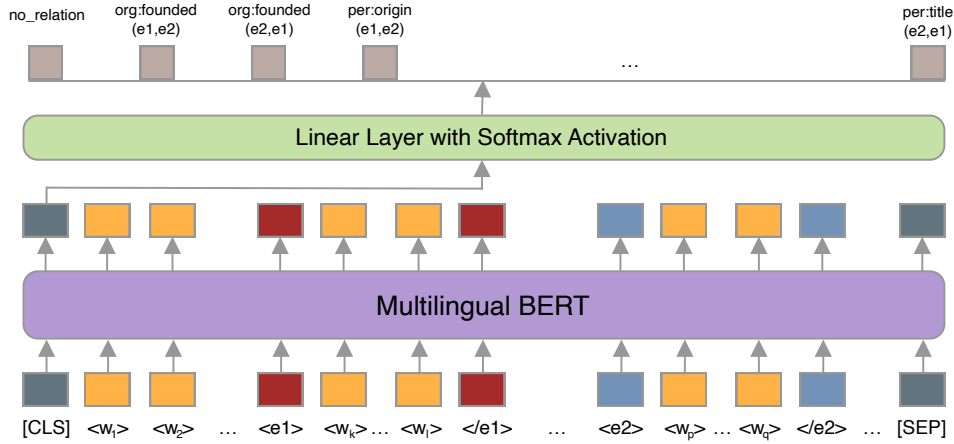


Figure 2: Illustration of our model. $\langle w_i \rangle$'s represent tokens from BERT tokenizer, $\langle e1 \rangle$, $\langle /e1 \rangle$ and $\langle e2 \rangle$, $\langle /e2 \rangle$ represent entity start and end markers for the first and second entities, respectively. [CLS] and [SEP] are special tokens in BERT. [CLS] can be used as a fixed-length input representation and [SEP] denotes the end of the sentence.

entity pairs.

$$\begin{aligned}
 D_s &= \{(S_i^s, E1_i^s, E2_i^s, r_i)\}_{i=1}^{n_s} \text{ where} \\
 S_i^s &= [w_1, w_2, \dots, w_n] \\
 E1_i^s &= (w_k, w_{k+1}, \dots, w_l) \\
 E2_i^s &= (w_p, w_{p+1}, \dots, w_q) \\
 r_i &\in R
 \end{aligned}$$

$E1_i^s$ and $E2_i^s$ correspond to entities and w_i correspond to tokens in the sentence S_i^s . r_i is the directional relation between $E1_i^s$ and $E2_i^s$ in S_i^s , selected from a predefined relation set R .

Given test set $D_t = \{(S_j^t, E1_j^t, E2_j^t)\}_{j=1}^{n_t}$ in the target language, cross-lingual relation classification aims to find the relation probability $P(r_j | S_j^t, E1_j^t, E2_j^t)$ where $r_j \in R$ for a sentence and an entity pair in the target language with the supervision of D_s in the source language.

4.1 Multilingual BERT

Multilingual BERT (mBERT) (Devlin et al., 2019), is a multilingual language model trained on 104 languages using the corresponding Wikipedia dumps. Due to shared word pieces like URLs and numbers across languages (Pires et al., 2019), mBERT is able to produce fixed-length sentence representations for these languages. Exponential smoothed weighting is used in order to reduce the under-representation problem of low-resource languages that have a relatively smaller number of Wikipedia articles.

mBERT is selected as our baseline model in this work, similar to recent cross-lingual tasks such as

natural language inference (Wu and Dredze, 2019) and question answering (Artetxe et al., 2020). Each sentence is tokenized by the mBERT tokenizer. Following (Soares et al., 2019), entity markers are added to emphasize the locations of the entities in the sentences. We add entity start and end markers that are special tokens, which are learned from scratch during the training, as shown in Figure 2.

Our objective is to predict the relation between a given entity pair in a sentence from among a set of relations. For this purpose, as in (Devlin et al., 2019), mBERT's output state of the [CLS] token is used as fixed-length sentence representation (or in our case as relation representation). This representation is fed into a linear layer with softmax activation to predict the probability of each relation, as illustrated in Figure 2. The developed model predicts the probabilities of the *no_relation* class and 18 directional relation classes, which result in 37 different classes in the KBP-37 and RELX datasets.

Our implementation details about mBERT are as follows.

- We use the initial weights of Cased Multilingual BERT from (Devlin et al., 2019), which has 12 layers, 768 hidden size, 12 heads, and 110M parameters.
- The network on top of the transformer architecture that gets the [CLS] representation as input for relation class prediction has a linear layer with softmax activation.
- AdamW with $3e - 5$ learning rate and 0.1 weight decay is used with a batch size of 64.

Sentences	
S_{en}	In the 3rd century, E2 wrote his “ E1 ” and other exegetical and theological works while living in Caesarea.
S_{es}	Este es un palimpsesto de una copia de la obra de E2 llamada la E1 . ³
S_{tr}	İreneyus ve E2 gibi kilise babalarının metinlerinde aktarılanlara göre esasen E3 li olan Marcellina, Anicetus döneminde Roma’ya göç etmiş ve çok sayıda takipçi toplamıştır. ⁴
Entities	
E1	<i>Q839739</i> (Hexapla, Hexapla, Hexapla)
E2	<i>Q170472</i> (Origen, Orígenes, Origenes)
E3	<i>Q87</i> (Alexandria, Alejandría, İskenderiye)
Relations	
(E1, E2)	<i>P50</i> (Author)
(E2, E3)	<i>P19</i> (Place of Birth)
Pairs	
Positive	(S_{en}, S_{es})
Negative	(S_{en}, S_{tr})

Figure 3: Sample positive and negative pairs constructed from RELX-Distant. Entities and relations are linked with their Wikidata ID’s (shown in italic) and words in parentheses in entities represent English, Spanish, and Turkish correspondence, consecutively.

- The classification loss is selected as the cross-entropy of the predictions with respect to the true labels.

The hyperparameters are tuned over the KBP-37 validation set based on the F1-score as described in Section 5. Learning rates of 1e-3, 1e-4, 3e-4, 1e-5, 3e-5, and 1e-6; weight decay values of 0, 0.01, and 0.1 with the SGD, Adam (Kingma and Ba, 2015), and AdamW (Loshchilov and Hutter, 2018) optimizers have been evaluated with PyTorch (Paszke et al., 2019) and HuggingFace’s Transformers (Wolf et al., 2019) libraries. The best values have been determined as $3e - 5$ learning rate and 0.1 weight decay with the AdamW optimizer.

4.2 Matching the Multilingual Blanks

Our objective is to pretrain a public checkpoint of mBERT, released by (Devlin et al., 2019), in a way that it can learn various representations of relations across different languages. In order to do this, we prepare RELX-Distant, whose entities are labeled by using Wikipedia hyperlinks, to create pairs of sentences from different languages and propose

³English Translation: This is a palimpsest of a copy of **E2**’s work called **E1**.

⁴English Translation: According to what was reported in the texts of the church fathers such as Irenaeus and **E2**, Marcellina, who was originally from **E3**, migrated to Rome during the Anicetus period and collected many followers.

Matching the Multilingual Blanks, a multilingual distant supervision approach that targets detecting the similarity between the relations described in an input multilingual pair of sentences.

For this model, we pretrain mBERT with two objectives: Masked Language Model from (Devlin et al., 2019) and Matching the Multilingual Blanks (MTMB). Similar to the monolingual work in (Soares et al., 2019), we create positive and negative multilingual sentence pairs from RELX-Distant for the MTMB objective. We pretrain mBERT with the aim of learning how relations are represented in different languages by predicting whether the English sentence and the non-English sentence in a pair have the same relation or not.

Positive sentence pairs are selected to share the same entities, which result in having the same relation by the distant supervision scheme. (S_{en}, S_{es}) in Figure 3 is a positive pair because both sentences include the E1 (Hexapla) and E2 (Origen) entities that have the P50 (Author) relation.

In the negative sentence pairs, each sentence has entities with different relations. In order to avoid dissimilar sentences in a negative pair, which may cause our model to make predictions based on the topics of the sentences, we use *strong* negative pairs similar to (Soares et al., 2019). In strong negative pairs, one of the entities in each sentence in the pair is common. (S_{en}, S_{tr}) in Figure 3 is a

strong negative pair because both sentences share the entity E2 (Origen), but the English sentence has the P50 (Author) relation, and the Turkish sentence has the P19 (Place of Birth) relation.

In the compiled sentence pairs, the entities are replaced by a special [BLANK] token with 0.7 probability to capture the text patterns better and avoid memorizing the entities. By following these steps, we create 20 million pairs of sentences from RELX-Distant to pretrain mBERT. These sentence pairs have a uniform distribution with respect to the positive and negative classes as well as the languages in RELX-Distant. We call the pretraining procedure of mBERT with multilingual sentence pairs, Matching the Multilingual Blanks (MTMB).

The implementation details of the model are similar to the model described in Section 4.1. However, before multi-way relation classification training, we first pretrain the public checkpoint of mBERT (Devlin et al., 2019) with two objectives. The first objective is the Masked Language Model, and we implement it as implemented in (Devlin et al., 2019). The second objective is a binary classification of sentence pairs, whether two sentences in different languages have the same relation or not. While fine-tuning mBERT in Section 4.1 is relatively inexpensive (less than 10 minutes in each epoch on a GPU), one epoch of MTMB with 20 millions of sentence pairs takes approximately 10 days on a Tesla V100 GPU. Considering this, we release the weights of our MTMB model publicly in <https://github.com/boun-tabii/RELX>.

5 Results

We compare our monolingual relation classification results using KBP-37 and the cross-lingual results using RELX. We report our results by taking the average scores of 10 runs to decrease the effect of high variance between different runs in BERT as stated in (Dodge et al., 2020).

Evaluation Metric: We use (18+1)-way evaluation by taking directionality into account as used in (Hendrickx et al., 2010). First, the F1 score of a relation is calculated by taking the micro average of F1’s of both directions. Then, the macro average of F1 scores of 18 relations is considered as our final score.

5.1 KBP-37

Table 3 contains the results for our models and the state-of-the-art models evaluated on the KBP-37

Model	Dev	Test
BERT _{Large} (Soares et al., 2019)	69.5	68.3
MTB (Soares et al., 2019)	70.3	69.3
BERT _{Base}	66.0	65.4
mBERT	65.5	64.9
MTMB	66.8	66.5

Table 3: F1 scores of our models compared to the state-of-the-art models on the development and test sets of KBP-37 (English).

Model	EN	FR	DE	ES	TR
mBERT	61.8	58.3	57.5	57.9	55.8
MTMB	63.6	59.9	59.9	62.4	56.2

Table 4: F1 scores of mBERT and MTMB evaluated on RELX. The columns represent the English, French, German, Spanish, and Turkish parts of RELX.

development and test sets. BERT_{Large} and MTB are models from (Soares et al., 2019). Both models use pretrained BERT_{Large}, which is specific to the English language. We finetune three models for relation classification with the same architecture and number of parameters: BERT_{Base}, mBERT, and MTMB; where mBERT and MTMB are pretrained on multilingual corpora, while BERT_{Base} is pretrained on English corpora. The complexity of BERT_{Large} is much higher than mBERT and BERT_{Base}. The number of parameters in BERT_{Large} is 340 million, while mBERT and BERT_{Base} have 110 million parameters. Also, BERT_{Large} has 24 layers and 16 heads compared to 12 layers and 12 heads in mBERT and BERT_{Base}. Finally, the hidden size in BERT_{Large} is 1024, while it is 768 in mBERT and BERT_{Base}. Because of the difference in complexity and the language of the training data, as expected, BERT_{Large} based models have better performance for the English language than mBERT based models. Still, the results show that Matching the Multilingual Blanks significantly (p-value < 0.05) outperforms mBERT and BERT_{Base} in the English language according to the randomization tests (Yeh, 2000).

5.2 RELX

RELX is used to evaluate the mBERT and MTMB models, which are finetuned on the training set (which is in English) of KBP-37. The results are summarized in Table 4. By Matching the Multilingual Blanks training setup, we significantly (p-value < 0.05) improve the performance

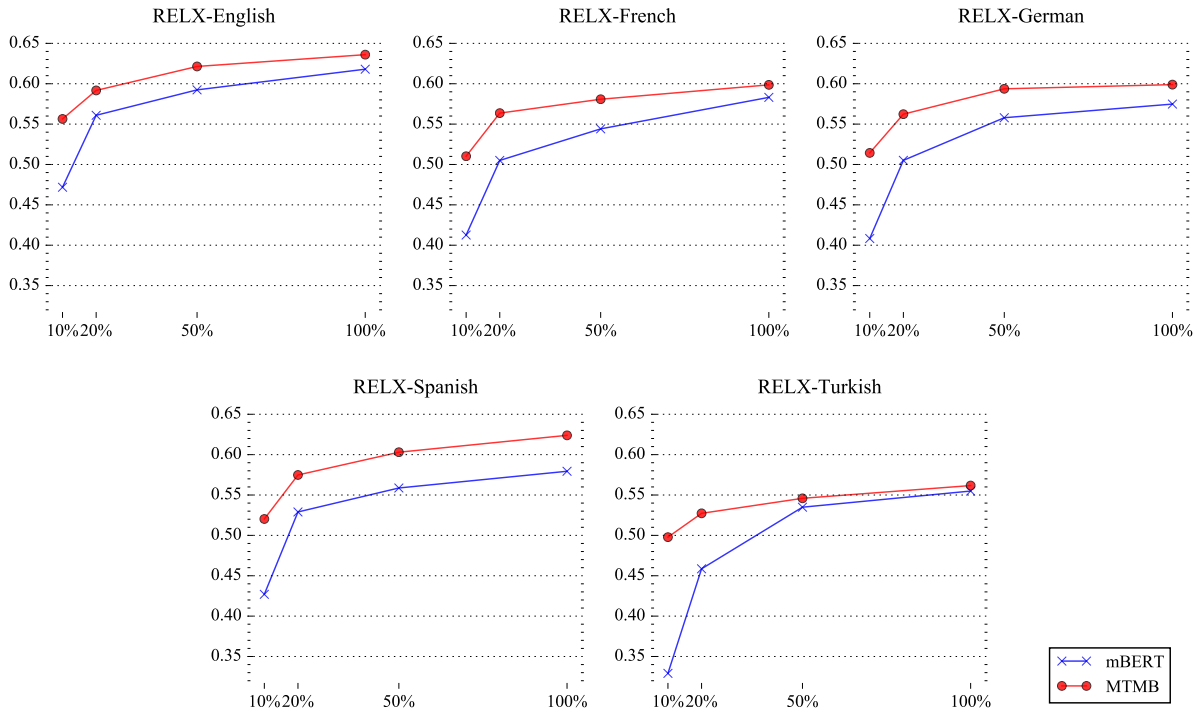


Figure 4: Cross-lingual relation classification performance (F1 score *y-axis*) of mBERT and MTMB with varying amounts of training data (*x-axis*).

of mBERT for five languages, including RELX-English. In cross-lingual cases and the monolingual case, MTMB significantly outperforms mBERT based on the randomization tests.

We display the results by varying the size of the training data in Figure 4. The results show that MTMB performs better than mBERT, especially in low-resource cases. The difference in F1 scores between MTMB and mBERT is more significant when the amount of the available training data is lower. For Spanish, MTMB was able to reach the performance of mBERT that uses all the training data by using only around 20% of the training data, and for the other evaluated languages (except Turkish), around 50% of the data was sufficient to obtain the same performance as mBERT that uses all the training data. Thus, the required human annotations in the source language can be significantly reduced with the help of MTMB.

Table 4 demonstrates that the best cross-lingual performance is achieved for Spanish, which is on par with prior studies on other cross-lingual NLP tasks such as question answering and natural language inference (Artetxe et al., 2020) that also report higher performance for Spanish. On the other hand, our results show that the worse cross-

lingual performance is obtained for Turkish. Pires et al. (2019) observe that mBERT performance is effected by word ordering and works best for typologically similar languages. In order to investigate this, we compare the source language (English) and target languages (French, German, Spanish, Turkish) by a subset of the World Atlas of Language Structures (WALS) features (Dryer and Haspelmath, 2013) that are relevant to grammatical ordering⁵ as in (Pires et al., 2019). Considering these features, Turkish is the least similar language to English among the languages in RELX. Our results support the claim presented in (Pires et al., 2019).

Error analysis reveals that 120 out of 176 mispredicted sentences in RELX-English are common in all target languages. Among these common errors, classes with less than 600 samples in the training data have 60% more error rate, suggesting that increasing their number of samples may benefit in all languages.

We also analyzed relation direction errors, where the predicted relation class is the same as the gold class, while the predicted direction is incorrect.

⁵81A: Order of Subject, Object and Verb, 85A: Order of Adposition and Noun Phrase, 86A: Order of Genitive and Noun, 87A: Order of Adjective and Noun, 88A: Order of Demonstrative and Noun, 89A: Order of Numeral and Noun

There are 79 relation direction errors for Turkish, whereas there are less than 15 for the other languages. Turkish has generally an SOV word order and postpositions, while English has generally SVO word order and prepositions. These differences between Turkish and English are possible causes for the problems related to direction errors as discussed in (Pires et al., 2019). Finally, no notable difference is observed in errors across languages in terms of sentence length.

6 Conclusion

In this paper, we addressed the cross-lingual relation classification task. First, we introduced two publicly available datasets: **RELX**, a cross-lingual relation classification benchmark for English, French, German, Spanish, and Turkish with parallel sentences and **RELX-Distant**, a multilingual dataset containing a large number of sentences with relations from Wikipedia and Wikidata collected via distant supervision. Second, we proposed a baseline model with mBERT and a new multilingual pretraining scheme with distant supervision called Matching the Multilingual Blanks (MTMB). Our experiments showed that MTMB significantly outperforms the mBERT baseline on the monolingual and cross-lingual datasets. The improvement obtained by MTMB is higher in the low-resource settings for the source language. We also showed that better cross-lingual relation classification performance is obtained for target languages which are typologically similar to the source language. The performance for Spanish is comparable to English (the source language in this study), while the lowest F1 scores are obtained for Turkish. MTMB can be easily adopted to other languages by using our provided scripts⁶. The only requirement is the availability of Wikipedia articles in the new target language.

As future work, we plan to extend RELX-Distant to all the available languages in Wikipedia. We will also investigate the effect of MTMB in different cross-lingual tasks such as natural language inference, named entity recognition, and question answering by using the extended RELX-Distant dataset.

Acknowledgments

We would like to thank Ahmed Yusuf Asiltürk, Barbaros Eriş, Esmanur Yılmaz, Gizem Özkanal,

⁶<https://github.com/boun-tabii/RELX>

Ramazan Pala, and Taha Küçükkatırcı for their contributions on the translation and eltur.co for their quality check service. TUBITAK-BIDEB 2211-A Scholarship Program (to A.K.), and TUBA-GEBIP Award of the Turkish Science Academy (to A.O.) are gratefully acknowledged.

References

- Mostafa Abdou, Cezar Sas, Rahul Aralikkatte, Isabelle Augenstein, and Anders Søgaard. 2019. X-WikiRE: A large, multilingual resource for relation extraction as machine comprehension. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 265–274.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. On the cross-lingual transferability of monolingual representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online. 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. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems*, pages 7059–7069.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith. 2020. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *arXiv preprint arXiv:2002.06305*.

- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Manaal Faruqui and Shankar Kumar. 2015. [Multilingual open relation extraction using cross-lingual projection](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1351–1356, Denver, Colorado. Association for Computational Linguistics.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. [SemEval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 33–38, Uppsala, Sweden. Association for Computational Linguistics.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Nitin Indurkha. 2015. [Emerging directions in predictive text mining](#). *WIREs Data Mining and Knowledge Discovery*, 5(4):155–164.
- Nanda Kambhatla. 2004. [Combining lexical, syntactic, and semantic features with maximum entropy models for information extraction](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 178–181, Barcelona, Spain. Association for Computational Linguistics.
- Seokhwan Kim, Minwoo Jeong, Jonghoon Lee, and Gary Geunbae Lee. 2010. [A cross-lingual annotation projection approach for relation detection](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 564–571, Beijing, China. Coling 2010 Organizing Committee.
- Seokhwan Kim and Gary Geunbae Lee. 2012. [A graph-based cross-lingual projection approach for weakly supervised relation extraction](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 48–53, Jeju Island, Korea. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Jiahua Liu, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2019. [XQA: A cross-lingual open-domain question answering dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2358–2368, Florence, Italy. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2018. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems*, pages 3111–3119.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011, Suntec, Singapore. Association for Computational Linguistics.
- Thien Huu Nguyen and Ralph Grishman. 2015. [Relation extraction: Perspective from convolutional neural networks](#). In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 39–48, Denver, Colorado. Association for Computational Linguistics.
- Jian Ni and Radu Florian. 2019. [Neural cross-lingual relation extraction based on bilingual word embedding mapping](#). 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 399–409, Hong Kong, China. Association for Computational Linguistics.
- Joel Nothman, Nicky Ringland, Will Radford, Tara Murphy, and James R Curran. 2013. [Learning multilingual named entity recognition from Wikipedia](#). *Artificial Intelligence*, 194:151–175.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems*, pages 8026–8037.
- 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.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. Matching the blanks: Distributional similarity for relation learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2895–2905.
- Shyam Upadhyay. 2019. *Exploiting Cross-lingual Representations for Natural Language Processing*. Ph.D. thesis, University of Pennsylvania.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. [ACE 2005 multilingual training corpus](#). Philadelphia: Linguistic Data Consortium.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- 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.
- Jiateng Xie, Zhilin Yang, Graham Neubig, Noah A. Smith, and Jaime Carbonell. 2018. [Neural cross-lingual named entity recognition with minimal resources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 369–379, Brussels, Belgium. Association for Computational Linguistics.
- Kun Xu, Siva Reddy, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2016. [Question answering on Freebase via relation extraction and textual evidence](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2326–2336, Berlin, Germany. Association for Computational Linguistics.
- Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015. Classifying relations via long short term memory networks along shortest dependency paths. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 1785–1794.
- Alexander Yeh. 2000. [More accurate tests for the statistical significance of result differences](#). In *COLING 2000 Volume 2: The 18th International Conference on Computational Linguistics*.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. [Relation classification via convolutional deep neural network](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Dongxu Zhang and Dong Wang. 2015. Relation classification via recurrent neural network. *arXiv preprint arXiv:1508.01006*.
- Bowei Zou, Zengzhuang Xu, Yu Hong, and Guodong Zhou. 2018. [Adversarial feature adaptation for cross-lingual relation classification](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 437–448, Santa Fe, New Mexico, USA. Association for Computational Linguistics.