

Contrastive Language Adaptation for Cross-Lingual Stance Detection

Mitra Mohtarami¹, James Glass¹, Preslav Nakov²

¹MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, USA

²Qatar Computing Research Institute, HBKU, Doha, Qatar

{mitra,glass}@csail.mit.edu; pnakov@hbku.edu.qa

Abstract

We study cross-lingual stance detection, which aims to leverage labeled data in one language to identify the relative perspective (or stance) of a given document with respect to a claim in a different target language. In particular, we introduce a novel contrastive language adaptation approach applied to memory networks, which ensures accurate alignment of stances in the source and target languages, and can effectively deal with the challenge of limited labeled data in the target language. The evaluation results on public benchmark datasets and comparison against current state-of-the-art approaches demonstrate the effectiveness of our approach.

1 Introduction

The rise of social media has enabled the phenomenon of “fake news,” which could target specific individuals and can be used for deceptive purposes (Lazer et al., 2018; Vosoughi et al., 2018). As manual fact-checking is a time-consuming and tedious process, computational approaches have been proposed as a possible alternative (Popat et al., 2017; Wang, 2017; Mihaylova et al., 2018, 2019), based on information sources such as social media (Ma et al., 2017), Wikipedia (Thorne et al., 2018), and knowledge bases (Huynh and Papotti, 2018). Fact-checking is a multi-step process (Vlachos and Riedel, 2014): (i) checking the reliability of media sources, (ii) retrieving potentially relevant documents from reliable sources as evidence for each target claim, (iii) predicting the stance of each document with respect to the target claim, and finally (iv) making a decision based on the stances from (iii) for all documents from (ii).

Here, we focus on stance detection which aims to identify the relative perspective of a document with respect to a claim, typically modeled using labels such as *agree*, *disagree*, *discuss*, and *unrelated*.

Current approaches to stance detection (Bar-Haim et al., 2017; Dungs et al., 2018; Kochkina et al., 2018; Sobhani et al., 2017; Mohtarami et al., 2018) are well-studied in mono-lingual settings, in particular for English, but less attention has been paid to other languages and *cross-lingual* settings. This is partially due to domain differences and to the lack of training data in other languages.

We aim to bridge this gap by proposing a cross-lingual model for stance detection. Our model leverages resources of a source language (e.g., English) to train a model for a target language (e.g., Arabic). Furthermore, we propose a novel contrastive language adaptation approach that effectively aligns samples with similar or dissimilar stances across source and target languages using task-specific loss functions. We apply our language adaptation approach to memory networks (Sukhbaatar et al., 2015), which have been found effective for mono-lingual stance detection (Mohtarami et al., 2018).

Our model can explain its predictions about stances of documents against claims in a different/target language by extracting relevant text snippets from the documents of the target language as evidence. We use evidence extraction as a measure to evaluate the transferability of our model. This is because more accurate evidence extraction indicates that the model can better learn semantic relations between claims and pieces of evidence, and consequently can better transfer knowledge to the target language.

The contributions of this paper are summarized as follows:

- We propose a novel language adaptation approach based on contrastive stance alignment that aligns the class labels between source and target languages for effective cross-lingual stance detection.

- Our model is able to extract accurate text snippets as evidence to explain its predictions in the target language (results are in Section 4.2).
- To the best of our knowledge, this is the first work on cross-lingual stance detection.

We conducted our experiments on English (as source language) and Arabic (as target language). In particular, we used the Fake News Challenge dataset (Hanselowski et al., 2018) as source data and an Arabic benchmark dataset (Baly et al., 2018) as target data. The evaluation results have shown 2.7 and 4.0 absolute improvement in terms of macro-F1 and weighted accuracy for stance detection over the current state-of-the-art mono-lingual baseline, and 11.4, 14.9, 16.1, 12.9, and 13.1 points of absolute improvement in terms of precision at ranks 1–5 for extracting evidence snippets respectively. Furthermore, a key finding in our investigation is that, in contrast to other tasks (Devlin et al., 2019; Peters et al., 2018), pre-training with large amounts of source data is less effective for cross-lingual stance detection. We show that this is because pre-training can considerably bias the model toward the source language.

2 Method

Assume that we are given a training dataset for a source language, \mathcal{D}_s , which contains a set of triplets as follows: $\mathcal{D}_s = \{((c_i^s, d_i^s), y_i^s)\}_{i=1}^N$, where N is the number of source samples, (c_i^s, d_i^s) is a pair of claim c_i^s and document d_i^s , and $y_i^s \in Y$, $Y = \{agree, disagree, discuss, unrelated\}$, is the corresponding label indicating the stance of the document with respect to the claim. In addition, we are given a very small training dataset for the target language $\mathcal{D}_t = \{((c_i^t, d_i^t), y_i^t)\}_{i=1}^M$, where M ($M \ll N$) is the number of target samples, and (c_i^t, d_i^t) is a pair of claim and document in the target language with stance label $y_i^t \in Y$. In reality, (i) the size of the target dataset is very small, (ii) claims and documents in the source and target languages are from different domains, and (iii) the only commonality between the source and target datasets is in their stance labels, i.e., $y_i^s, y_i^t \in Y$.

We develop a language adaptation approach to effectively use the commonality between the source and the target datasets in their label space and to deal with the limited size of the target training data. We apply our language adaptation approach to end-to-end memory networks (Sukhbaatar et al., 2015) for cross-lingual stance detection.

We use memory networks as they have achieved state-of-the-art performance for mono-lingual stance detection (Mohtarami et al., 2018). However, our language adaptation approach can be applied to any other type of neural network. The architecture of our cross-lingual stance detection model is shown in Figure 1. It has two main components: (i) *Memory Networks* indicated with two dashed boxes for the source and the target languages, and (ii) *Contrastive Language Adaptation* component. In what follows, we first explain our memory network model for cross-lingual stance detection (Section 2.1) and then present our contrastive language adaptation approach (Section 2.2).

2.1 Memory Networks

Memory networks are designed to remember past information (Sukhbaatar et al., 2015) and have been successfully applied to NLP tasks ranging from dialog (Bordes et al., 2017) to question answering (Xiong et al., 2016) and mono-lingual stance detection (Mohtarami et al., 2018). They include components that can potentially use different learning models and inference strategies. Our source and target memory networks follow the same architecture as depicted in Figure 1:

A memory network consists of six components. The network takes as input a document d and a claim c and encodes them into the *input space* I . These representations are stored in the *memory* component M for future processing. The relevant parts of the input are identified in the *inference* component F , and used by the *generalization* component G to update the memory M . Finally, the *output* component O generates an output from the updated memory, and encodes it to a desired format in the *response* component R using a prediction function, e.g., `softmax` for classification tasks. We elaborate on these components below.

Input representation component I : It encodes documents and claims into corresponding representations. Each document d is divided into a sequence of paragraphs $X = (x_1, \dots, x_l)$, where each x_j is encoded as \mathbf{m}_j using an LSTM network, and as \mathbf{n}_j using a CNN; these representations are stored in the memory component M . Note that while LSTMs are designed to capture and memorize their inputs (Tan et al., 2016), CNNs emphasize the local interaction between individual words in sequences, which is important for obtaining good representation (Kim, 2014).




Figure 1: The architecture of our cross-lingual memory network for stance detection.

Thus, our I component uses both LSTM and CNN representations. It also uses separate LSTM and CNN with their own parameters to represent each input claim c as \mathbf{c}_{lstm} and \mathbf{c}_{cnn} , respectively.

We consider each paragraph as a single piece of *evidence* because a paragraph usually represents a coherent argument, unified under one or more inter-related topics. We thus use the terms paragraph and evidence interchangeably.

Inference component F : Our inference component computes LSTM- and CNN-based similarity between each claim c and evidence x_j as follows:

$$P_j^{lstm} = \mathbf{c}_{lstm}^\top \times \mathbf{M} \times \mathbf{m}_j, \forall j$$

$$P_j^{cnn} = \mathbf{c}_{cnn}^\top \times \mathbf{M}' \times \mathbf{n}_j, \forall j$$

where P_j^{lstm} and P_j^{cnn} indicate claim-evidence similarity based on LSTM and CNN respectively, $\mathbf{c}_{lstm} \in \mathbb{R}^q$ and $\mathbf{m}_j \in \mathbb{R}^d$ are LSTM representations of c and x_j respectively, $\mathbf{c}_{cnn} \in \mathbb{R}^{q'}$ and $\mathbf{n}_j \in \mathbb{R}^{d'}$ are the corresponding CNN representations, and $\mathbf{M} \in \mathbb{R}^{q \times d}$ and $\mathbf{M}' \in \mathbb{R}^{q' \times d'}$ are similarity matrices trained to map claims and paragraphs into the same space with respect to their LSTM and CNN representations. The rationale behind using these similarity matrices is that, in the memory network, we seek a transformation of the input claim, i.e., $\mathbf{c}^\top \times \mathbf{M}$, in order to obtain the closest evidence to the claim.

Additionally, we compute another semantic similarity vector, P_j^{fidf} , by applying a cosine similarity between the TF.IDF (Spärck Jones, 2004) representation of x_j and c . This is particularly useful for stance detection as it can help filtering out unrelated pieces of evidence.

Memory M and Generalization G components:

Our memory component stores representations and the generalization component improves their quality by filtering out unrelated evidence. For example, the LSTM representations of paragraphs, $\mathbf{m}_j, \forall j$, are updated using the claim-evidence similarity P_j^{fidf} as follows: $\mathbf{m}_j = \mathbf{m}_j \odot P_j^{fidf}, \forall j$. This transformation will help filter out unrelated evidence with respect to claims. The updated \mathbf{m}_j in conjunction with \mathbf{c}_{lstm} are used by the inference component F to compute $P_j^{lstm}, \forall j$ as we explained above. Then, P_j^{lstm} are in turn used to update CNN representations in memory as follows: $\mathbf{n}_j = \mathbf{n}_j \odot P_j^{lstm}, \forall j$. Finally, the updated \mathbf{n}_j and \mathbf{c}_{cnn} are used to compute P_j^{cnn} .

Output representation component O : This component computes the output of the memory M by concatenating the average vector of the updated \mathbf{n}_j with the *maximum* and *average* of claim-evidence similarity vectors P_j^{fidf}, P_j^{lstm} and P_j^{cnn} . The maximum helps to identify parts of documents that are most similar to claims, while the average estimates the overall document-claim similarity.

Response generation component R : This component computes the final stance of a document with respect to a claim. For this, the output of component O is concatenated with \mathbf{c}_{lstm} and \mathbf{c}_{cnn} and fed into a `softmax` to predict the stance of the document with respect to the claim.

All the memory network parameters, including those of CNN and LSTM in the I component, the similarity matrices \mathbf{M} and \mathbf{M}' in F , and the classifier parameters in R , are jointly learned during the training process with our language adaptation.

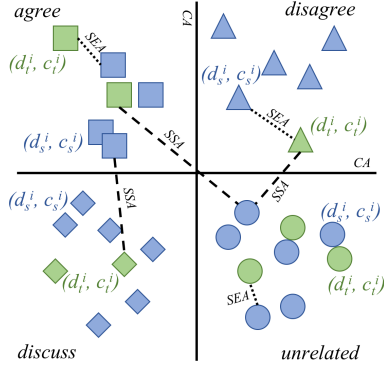


Figure 2: Illustration of stance equal alignment (SEA), stance separation alignment (SSA), and classification alignment (CA) constraints. Different shapes indicate different stance labels and colors specify source (*blue*) and target (*green*) languages.

2.2 Contrastive Language Adaptation

Memory networks are effective for stance detection in mono-lingual settings (Mohtarami et al., 2018) when there is sufficient training data. However, we show that these models have limited transferability to target languages with limited data. This could be due to discrepancy between the underlying data distributions in the source and target languages. We show that the performance of these networks can be trivially increased when the model, pre-trained on source data, is fine-tuned using small amounts of labeled target data. We further develop contrastive language adaptation that can exploit the labeled source data to perform well on target data.

Our contrastive adaptation approach:

- encourages pairs (d_i^s, c_i^s) from the source language and (d_i^t, c_i^t) from the target language with the same stance labels (i.e., $y_i^s = y_i^t$) to be nearby in the embedding space. We call this mapping *Stance Equal Alignment* (SEA), illustrated with dotted lines in Figure 2. Note that documents and claims in the two languages are often semantically different and are not corresponding translations of each other.
- encourages pairs (d_i^s, c_i^s) from the source language and (d_i^t, c_i^t) from the target language with different stance labels (i.e., $y_i^s \neq y_i^t$) to be far apart in the embedding space. We call this mapping *Stance Separation Alignment* (SSA), shown with dashed lines in Figure 2.
- encourages pairs (d_i^s, c_i^s) from the source language and (d_i^t, c_i^t) from the target language to be correctly classified as $(d_i^s, c_i^s) \rightarrow y_i^s$ and $(d_i^t, c_i^t) \rightarrow y_i^t$. We call this *Classification Alignment* (CA), solid lines in Figure 2.

Algorithm 1. Cross-Lingual Stance Detection Model

Inputs:

- $\{(d_i^s, c_i^s), y_i^s\}_1^{l^s}$: set of pairs of documents d_i^s and claims c_i^s in the source language, where $y_i^s \in Y$, $Y = \{agree, disagree, discuss, unrelated\}$
- $\{(d_i^t, c_i^t), y_i^t\}_1^{l^t}$: small set of pairs of documents d_i^t and claims c_i^t in the target language with $y_i^t \in Y$

Output:

- $\{(d_i^t, c_i^t) \rightarrow y_i^t\}_{i=1}^M$: assign stance labels $y_i^t \in Y$ to given unlabeled target pairs

Cross-Lingual model:

- 1 Create the sets $\{(d_i^s, c_i^s), (d_j^t, c_j^t), (y_i^s, y_i^t)\}$ and $\{(d_j^t, c_j^t), (d_i^s, c_i^s), (y_j^t, y_j^t)\}$, where $y' = 1$ if the source and the target have the same label, and $y' = 0$ otherwise.
- 2 **Loop for e epochs:**
- 3 pass (d_i^s, c_i^s) to the source memory network to create its representation r_i^s .
- 4 pass (d_j^t, c_j^t) to the target memory network to create its representation r_j^t .
- 5 pass (r_i^s, y_i^s) to the classification to compute its classification loss $L_{CA_i^s}$.
- 6 pass (r_i^s, r_i^t, y_i^t) to the language adaptation to compute the stance alignment loss L_{CSA_i} .
- 7 compute total loss $L_i^s = (1-\alpha)L_{CA_i^s} + \alpha L_{CSA_i}$.
- 8 repeat steps 2-7 with a change in step 5 by passing the target sample (r_j^t, y_j^t) to the classification instead of the source sample, and compute its L_j^t in step 7.
- 9 jointly optimize all parameters of the model using the average loss $L = \text{mean}(\{L_i^s\} + \{L_j^t\})$.

Table 1: Cross-lingual stance detection model.

We make complete use of the stance labels in the cross-lingual setting by parameterizing our model according to the distance between the source and the target samples in the embedding space.

For stance equal alignment (SEA) constraint, the objective is to minimize the distance between pairs of source and target data with the same stance labels. We achieve this using the following loss:

$$L_{SEA} = \begin{cases} \mathbb{D}(\mathfrak{g}(d_i^s, c_i^s) - \mathfrak{g}(d_i^t, c_i^t))^2, & y_i^s = y_i^t \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where \mathfrak{g} maps its input pair to an embedding space using our memory network or any mono-lingual model, and \mathbb{D} computes Euclidean distance.

For stance separation alignment (SSA), the goal is to maximize the distance between pairs with different stance labels. We use the following loss:

$$L_{SSA} = \begin{cases} \max(0, m - \mathbb{D}(\mathfrak{g}(d_i^s, c_i^s) - \mathfrak{g}(d_i^t, c_i^t))^2), & y_i^s \neq y_i^t \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

where we maximize the distance between pairs with different stance labels up to the margin m .

The margin parameter m specifies the extent of separability in the embedding space.

We can further use any classification loss to enforce classification alignment (CA). We use categorical cross-entropy and call it *Classification Alignment* loss L_{CA} .

We develop our overall language adaptation loss, named *Contrastive Stance Alignment* loss, L_{CSA} , by combining L_{SEA} and L_{SSA} as follows:

$$L_{CSA} = L_{SEA} + L_{SSA}. \quad (3)$$

Finally, the total loss of our cross-lingual stance detection model is defined as follows:

$$L = (1 - \alpha)L_{CA} + \alpha L_{CSA}, \quad (4)$$

where the α parameter controls the balance between classification and language adaptation losses, which we optimize on the validation dataset.

Information Flow: Our overall cross-lingual model for stance detection is shown in Figure 1, and a summary of the algorithm is presented in Table 1. As Figure 1 shows, each source and target pairs are passed to the source and to the target memory networks to obtain their corresponding representations (Lines 3-4 in Table 1). The source representation and its gold stance label are passed to the classifier to compute the classification loss (Line 5). In addition, the source and the target representations in conjunction with a binary parameter (y' , which is 1 if the source and the target have the same stance label, and 0 otherwise) are passed to the language adaptation component to compute the contrastive stance alignment loss L_{CSA} (Line 6). Finally, the total loss is computed based on Equation (4) (Line 7).

The classifier also uses labeled target samples to create a shared embedding space and to fine-tune itself with respect to the target language. For this purpose, we repeat the above steps by switching the target and the source pipelines (Line 8). Finally, we compute the average of all losses and we use it to optimize the parameters of our model (Line 9).

Pre-training for Language Adaptation: Pre-training has been found effective in many language adaptation settings (Tzeng et al., 2017). To investigate the effect of pre-training, we first pre-train the source memory network and the classifier using \mathcal{D}_s (only the top pipeline in Figure 1), and then we apply language adaptation with the full model.

3 Experiments

Data and Settings. As source data, we use the Fake News Challenge dataset¹ which contains 75.4K claim-document pairs in English with $\{agree, disagree, discuss, unrelated\}$ as stance labels. As target dataset, we use 3K Arabic claim-document pairs developed in (Baly et al., 2018).²

We perform 5-fold cross-validation on the Arabic dataset, using each fold in turn for testing, and keeping 80% of the remaining data for training and 20% for development. We use 300-dimensional pretrained cross-lingual Wikipedia word embeddings from MUSE (Lample et al., 2018). We use 300-dimensional units for the LSTM and 100 feature maps with filter width of 5 for the CNN. We consider the first 9 paragraphs per document, which is the median number of paragraphs in source documents. We optimize all hyper-parameters on validation data using *Adam* (Kingma and Ba, 2014).

Evaluation Measures. We use the followings:

- *Accuracy*: The fraction of correctly classified examples. For multi-class classification, accuracy is equivalent to micro-average F_1 (Manning et al., 2008).
- *Macro- F_1* : The average of F_1 scores that were calculated for each class separately.
- *Weighted Accuracy*: This is a hierarchical metric, which first awards 0.25 points if the model correctly predicts a document-claim pair as *related*³ or *unrelated*. If it is *related*, 0.75 additional points are assigned if the model correctly predicts the pair as *agree*, *disagree*, or *discuss*. The goal of this weighting schema is to balance out the large number of *unrelated* examples (Hanselowski et al., 2018).

Baselines. We consider the following baselines:

- **Heuristic**: Given the imbalanced nature of our data, we use two heuristic baselines where all test examples are labeled as *unrelated* or *agree*. The former is a majority class baseline favoring accuracy and macro- F_1 , while the latter is better for weighted accuracy.
- **Gradient Boosting** (Baly et al., 2018): This is a Gradient Boosting classifier with n -gram features as well as indicators for refutation and polarity.

¹Available at www.fakenewschallenge.org

²Available at <http://groups.csail.mit.edu/sls/downloads/>

³*related* = $\{agree, disagree, discuss\}$

Methods	Weigh. Acc.	Acc.	Macro-F1	F1 (<i>agree, disagree, discuss, unrelated</i>)
1. All-unrelated	34.8	68.1	20.3	0 / 0 / 0 / 81.0
2. All-agree	40.2	15.6	6.7	27.0 / 0 / 0 / 0
3. Gradient Boosting (Baly et al., 2018)	55.6	72.4	41.0	60.4 / 9.0 / 10.4 / 84.0
4. TFMLP (Riedel et al., 2017)	49.3	66.0	37.1	47.0 / 7.8 / 13.4 / 80.0
5. EnrichedMLP (Hanselowski et al., 2018)	55.1	70.5	41.3	59.1 / 9.2 / 14.1 / 82.3
6. Ensemble (Baird et al., 2017)	53.6	71.6	37.2	57.5 / 2.1 / 6.2 / 83.2
7. MN <i>target</i> → <i>target</i> (Mohtarami et al., 2018)	55.3	70.9	41.7	60.0 / 15.0 / 08.5 / 83.1
8. MN <i>source</i> → <i>target</i>	53.2	64.2	36.0	40.4 / 02.0 / 19.1 / 82.4
9. MN (<i>source, target</i>) → <i>target</i>	57.3	65.0	42.5	58.0 / 12.2 / 20.9 / 79.0
10. ADMN (<i>adversarial</i>)	58.6	58.4	43.4	60.2 / 16.1 / 24.2 / 72.9
11. CLMN (<i>contrastive</i>)	61.3	71.6	45.2	65.1 / 11.6 / 20.5 / 83.7

Table 2: Evaluation results on the target Arabic test dataset.

- **TFMLP** (Riedel et al., 2017): This is an MLP with normalized bag-of-words features enriched with a single TF.IDF-based similarity feature for each claim-document pair.
- **EnrichedMLP** (Hanselowski et al., 2018): This model combines five MLPs, each with six hidden layers and advanced features from topic models, latent semantic analysis, etc.
- **Ensemble** (Baird et al., 2017): It is an ensemble based on weighted average of a deep convolutional neural network and a gradient-boosted decision tree - the best model at the Fake News Challenge.
- **Mono-lingual Memory Network** (Mohtarami et al., 2018): This model is the current state-of-the-art for stance detection on our source dataset. It is an end-to-end memory network which incorporates both CNN and LSTM for prediction.
- **Adversarial Memory Network**: We use adversarial domain adaptation (Ganin et al., 2016) instead of contrastive language adaptation in our cross-lingual memory network.

Results. Table 2 shows the performance of all models on the target *Arabic* test set. The All-*unrelated* and All-*agree* baselines perform poorly across evaluation measures; All-*unrelated* performs better than All-*agree* because *unrelated* is the dominant class ($\sim 68\%$ of examples).

Rows 3–6 show that *Gradient Boosting* and *EnrichedMLP* yield similar results, while *TFMLP* performs the worst. We attribute this to the advanced features used in the two former models. *Gradient Boosting* has better accuracy due to its better performance on the dominant class. Note that *Ensemble* performs poorly because of the limited labeled data, which is insufficient to train a good CNN model.

Rows 7–9 show the results for the mono-lingual memory network (MN) from (Mohtarami et al., 2018). The performance of this model when trained on Arabic data only (row 7) is comparable to previous baselines (rows 3–6). But, it shows poor performance if trained on source English data and tested on Arabic test data (row 8). The model performs best (in terms of weighted accuracy and F1) if first pretrained on source data and then fine-tuned on target training data (row 9).

Row 10 in Table 2 shows the results for adversarial memory network (ADMN). It improves the performance of mono-lingual MN on weighted accuracy and F1, but its accuracy significantly drops. This is because adversarial approaches give higher weights to samples of the majority class (i.e., *unrelated*) which makes classification more challenging for the discriminator (Montahaei et al., 2018).

Row 11 shows the results for our cross-lingual memory network (CLMN) with ($\alpha = .7$); α controls the balance between classification and language adaptation losses (tuned using validation data). CLMN outperforms other baselines in terms of weighted accuracy and F1 while showing comparable accuracy. We show that the improvement is due to language adaptation being able to effectively transfer knowledge from the source to target language (see Section 4.2).

The last column in Table 2 shows that *unrelated* examples are the easiest ones. Also, although the *agree* and the *discuss* classes have roughly the same size, i.e., 474 and 409 examples, respectively, the results for *agree* are notably higher. This is mainly because the documents that discuss a claim often share the same topic with the claim, but they do not take a stance. In addition, the *disagree* examples are the most difficult ones; this class is by far the smallest one, with only 87 examples.

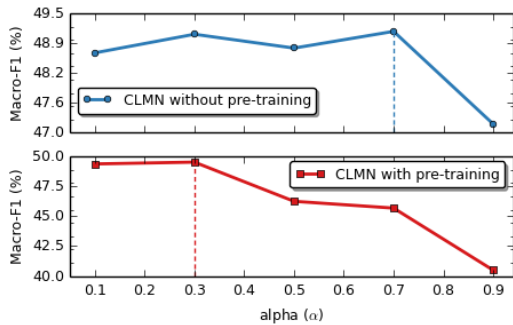


Figure 3: Impact of pretraining on the macro-F1 across α values. The Y axes show average results on the validation datasets with 5-fold cross-validation.

Methods	Weigh. Acc.	Acc.	Macro-F1
CLMN (with pretraining)	60.2	69.8	43.2
CLMN (without pretraining)	61.3	71.6	45.2

Table 3: CLMN results on the target test dataset.

4 Discussion

4.1 Effect of Pretraining

Table 3 shows CLMN without pretraining ($\alpha = .7$) performs better on target test data than CLMN with pretraining ($\alpha = .3$), recall that α controls the balance between classification and language adaptation losses. Our further analysis shows that pretraining biases the model toward the source language. Figure 3 shows the impact of pretraining on macro-average F1 score for CLMN across different values of α on validation data. While the model without pretraining achieves its best performance with a large α ($\alpha = 0.7$), the model with pretraining performs well with a smaller α ($\alpha = 0.3$). This suggests that our model can capture the characteristics of the source dataset via pretraining when using small supervision from language adaptation (i.e., small α). However, pretraining introduces bias to the source space and the performance drops when larger weights are given to language adaptation; see the results with pretraining in Figure 3.

4.2 Assessment of Model Transferability

The improvements of CLMN model over the mono-lingual MN models that use the target only, the source only, or both the target and the source (rows 7–9 in Table 2 respectively) indicate its transferability. We further estimate transferability by measuring the accuracy of the models in extracting evidence that support their predictions. A more accurate model should better transfer knowledge to the target language by accurately learning the relations between claims and pieces of evidence.

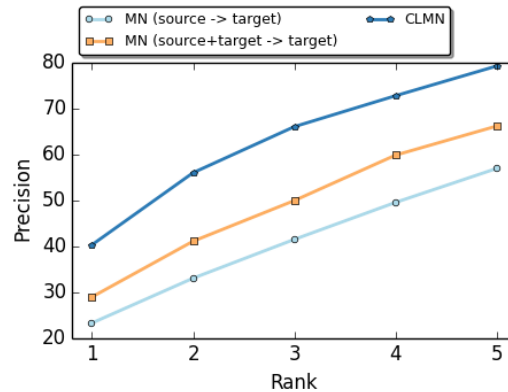


Figure 4: Transferability of our cross-lingual model.

Our target data has annotations (in terms of binary labels) for each piece of evidence (here paragraph) that indicate whether it is a rationale for the *agree* or for the *disagree* class. Moreover, our inference component (I) has a claim-evidence similarity vector, P_j^{cnn} , which can be used to rank pieces of evidence from the target document against the target claim.

We use the gold data and the rankings produced by our model in order to measure its precision in extracting evidence that supports its predictions. Figure 4 shows that our CLMN model achieves precision of 40.2, 55.9, 66.0, 72.7, and 79.2 at ranks 1–5 respectively, and outperforms mono-lingual MN models. This indicates that CLMN can better generalize and transfer knowledge to the target language through learning relations between pieces of evidence and claims.

4.3 Effect of Language Adaptation

Figures 5a and 5b show the classification (L_{CA}) versus contrastive stance (L_{CSA}) losses obtained from our best language adaptation model (i.e., without pretraining) across training epochs and α values. The results are averaged on validation data when performing 5-fold cross-validation. As Figure 5a shows, there is greater reduction in the classification loss for smaller values of α , i.e., when classification loss contributes more to the overall loss; see Equation (4). On the other hand, Figure 5b shows that the CSA loss decreases with larger values of α as the model pays more attention to the CSA loss; see the red and green lines in Figure 5b. These results indicate that our language adaptation model can find a good balance between the classification loss and the CSA loss, with the value of $\alpha = .7$ yielding the best performance.




Figure 5: Classification vs. contrastive stance alignment (CSA) losses across training epochs and α values.

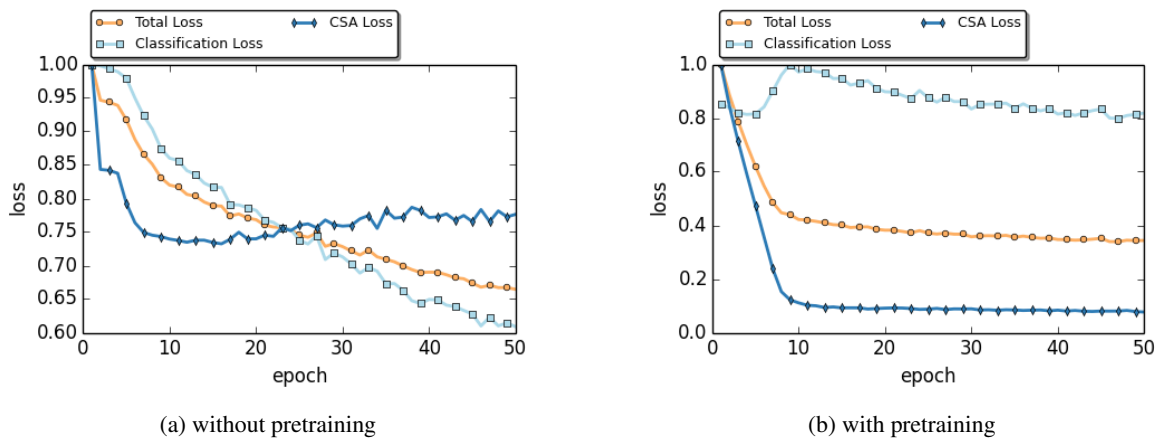


Figure 6: Classification loss vs. contrastive stance alignment loss (CSA) vs. total loss during training.

Figure 6 compares the classification L_{CA} , contrastive stance L_{CSA} , and total L losses obtained by our CLMN model on the validation dataset across training epochs when the loss weight parameter (α) is set to its best value. Figures 6a and 6b show the results without and with pretraining for $\alpha = .7$ and $\alpha = .3$ respectively. Without pretraining (Figure 6a), the classification (light-blue line) and CSA (dark-blue line) losses both decrease up to epoch 20, after which the classification loss keeps decreasing, but the CSA loss starts increasing. With pretraining (Figure 6b), the CSA loss rapidly decreases for the first 10 epochs (even though it has a small effect as $\alpha = .3$), and then continues with a smooth trend. This is because, during the initial training epochs, the model is biased to the *source* embedding space due to pretraining, and therefore the source and the target examples are far from each other. Then, our language adaptation model aligns the source and the target examples to form a much

better shared embedding space and this alignment strategy yields a rapid decrease of the CSA loss in the first few epochs. Yet, in contrast to the CSA loss, the classification loss increases in the first few epochs. This is because the model enforces alignment between the source and the target samples due to the large distances. Finally, the total loss (orange line) indicates a good balance between the classification and the language adaptation losses, and it consistently decreases during training.

5 Related Work

Domain Adaptation. Previous work has presented several domain adaptation techniques. Unsupervised domain adaptation approaches (Ganin and Lempitsky, 2015; Long et al., 2016; Muandet et al., 2013; Gong et al., 2012) attempt to align the distribution of features in the embedding space mapped from the source and the target domains.

A limitation of such approaches is that, even with perfect alignment, there is no guarantee that the same-label examples from different domains would map nearby in the embedding space. Supervised domain adaptation (Daumé and Marcu, 2006; Becker et al., 2013; Bergamo and Torresani, 2010) attempts to encourage same-label examples from different domains to map nearby in the embedding space. While supervised approaches perform better than unsupervised ones, recent work (Motiiian et al., 2017) has demonstrated superior performance by additionally encouraging class separation, meaning that examples from different domains and with different labels should be projected as far apart as possible in the embedding space. Here, we combined both types of alignments for cross-lingual stance detection.

Stance Detection. This is an important component for automatic fact-checking systems and veracity inference (Nadeem et al., 2019; Zhang et al., 2019; Atanasova et al., 2019). There have been some nuances in the way researchers have defined the stance detection task. Mohammad et al. (2016) and Zarrella and Marsh (2016) worked on stances regarding target propositions, e.g., entities, concepts or events, as *in-favor*, *against*, or *neither*. Most commonly, stance detection has been defined with respect to a *claim* as *agree*, *disagree*, *discuss* or *unrelated* (Hanselowski et al., 2018; Xu et al., 2018; Fang et al., 2019). Previous work mostly developed the models with rich hand-crafted features such as words, word embeddings, and sentiment lexicons (Riedel et al., 2017; Baird et al., 2017; Hanselowski et al., 2018). More recently, Mohtarami et al. (2018) presented a mono-lingual and feature-light memory network for stance detection. In this paper, we built on this work to extend previous efforts in stance detection to a cross-lingual setting, achieving the state-of-the-art result on the target language.

6 Conclusion and Future Work

We proposed an effective language adaptation approach to align class labels in source and target languages for accurate cross-lingual stance detection. Moreover, we investigated the behavior of our model in details and we have shown that it offers sizable performance gains over a number of competing approaches. In future, we will extend our language adaptation model to document retrieval and check-worthy claim detection tasks.

Acknowledgments

We thank the anonymous reviewers for their insightful comments. This research was supported in part by the Qatar Computing Research Institute, HBKU⁴ and DSTA of Singapore.

References

- Pepa Atanasova, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, Georgi Karadzhov, Tsvetomila Mihaylova, Mitra Mohtarami, and James Glass. 2019. Automatic fact-checking using context and discourse information. *ACM Journal of Data and Information Quality*, 11(3):12:1–12:27.
- Sean Baird, Doug Sibley, and Yuxi Pan. 2017. Talos targets disinformation with fake news challenge victory. <https://blog.talosintelligence.com/2017/06/talos-fake-news-challenge.html>.
- Ramy Baly, Mitra Mohtarami, James Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018. Integrating stance detection and fact checking in a unified corpus. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Volume 2*, pages 21–27, New Orleans, Louisiana, USA.
- Roy Bar-Haim, Indrajit Bhattacharya, Francesco Dinuzzo, Amrita Saha, and Noam Slonim. 2017. Stance classification of context-dependent claims. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, EACL ’17, pages 251–261, Valencia, Spain.
- Carlos J Becker, Christos M Christoudias, and Pascal Fua. 2013. Non-linear domain adaptation with boosting. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, pages 485–493. Lake Tahoe, Nevada.
- Alessandro Bergamo and Lorenzo Torresani. 2010. Exploiting weakly-labeled web images to improve object classification: a domain adaptation approach. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 1*, pages 181–189. Vancouver, British Columbia, Canada.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. Learning end-to-end goal-oriented dialog. In *Proceedings of the 5th International Conference on Learning Representations*, ICLR ’17, Toulon, France.

⁴This research is part of the Tanbih project (available at <http://tanbih.qcri.org/>) which aims to limit the effect of “fake news”, propaganda and media bias by making users aware of what they are reading.

- Hal Daumé, III and Daniel Marcu. 2006. Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research*, 26(1):101–126.
- 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 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*, pages 4171–4186, Minneapolis, Minnesota, USA.
- Sebastian Dungs, Ahmet Aker, Norbert Fuhr, and Kalina Bontcheva. 2018. Can rumour stance alone predict veracity? In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3360–3370, Santa Fe, NM, USA.
- Wei Fang, Moin Nadeem, Mitra Mohtarami, and James Glass. 2019. Neural multi-task learning for stance prediction. In *Proceedings of the EMNLP Workshop on Fact Extraction and Verification*, Hong Kong, China.
- Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on Machine Learning - Volume 37, ICML'15*, pages 1180–1189.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):2096–2030.
- Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. 2012. Geodesic flow kernel for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2066–2073.
- Andreas Hanselowski, Avinesh PVS, Benjamin Schiller, Felix Caspelherr, Debanjan Chaudhuri, Christian M. Meyer, and Iryna Gurevych. 2018. A retrospective analysis of the fake news challenge stance-detection task. In *Proceedings of the International Conference on Computational Linguistics, COLING '18*, pages 1859–1874, Santa Fe, NM, USA.
- Viet-Phi Huynh and Paolo Papotti. 2018. Towards a benchmark for fact checking with knowledge bases. In *Companion Proceedings of the The Web Conference 2018, WWW '18*, pages 1595–1598, Lyon, France.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP '14*, pages 1746–1751, Doha, Qatar.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Elena Kochkina, Maria Liakata, and Arkaitz Zubiaga. 2018. All-in-one: Multi-task learning for rumour verification. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3402–3413, Santa Fe, NM, USA.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *Proceedings of the 6th International Conference on Learning Representations*, Vancouver, BC, Canada.
- David M.J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science*, 359(6380):1094–1096.
- Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. 2016. Unsupervised domain adaptation with residual transfer networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 136–144, Barcelona, Spain.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL '17*, pages 708–717, Vancouver, Canada.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA.
- Tsvetomila Mihaylova, Georgi Karadzhov, Pepa Atanasova, Ramy Baly, Mitra Mohtarami, and Preslav Nakov. 2019. SemEval-2019 task 8: Fact checking in community question answering forums. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 860–869, Minneapolis, Minnesota, USA.
- Tsvetomila Mihaylova, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, Mitra Mohtarami, Georgi Karadzhov, and James Glass. 2018. Fact checking in community forums. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, New Orleans, LA, USA.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiao-Dan Zhu, and Colin Cherry. 2016. SemEval-2016 task 6: Detecting stance in tweets. In *Proceedings of Semantic Evaluation, SemEval '16*, pages 31–41, Berlin, Germany.

- Mitra Mohtarami, Ramy Baly, James Glass, Preslav Nakov, Lluís Màrquez, and Alessandro Moschitti. 2018. Automatic stance detection using end-to-end memory networks. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (Long Papers)*, pages 767–776, New Orleans, LA, USA.
- Ehsan Montahaei, Mahsa Ghorbani, Mahdiah Soleymani Baghshah, and Hamid R. Rabiee. 2018. Adversarial classifier for imbalanced problems. *CoRR*, abs/1811.08812.
- Saeid Motiian, Marco Piccirilli, Donald A. Adjeroh, and Gianfranco Doretto. 2017. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5715–5725, Venice, Italy.
- Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. 2013. Domain generalization via invariant feature representation. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, pages 10–18, Atlanta, GA, USA.
- Moin Nadeem, Wei Fang, Brian Xu, Mitra Mohtarami, and James Glass. 2019. FAKTA: An automatic end-to-end fact checking system. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 78–83, Minneapolis, MN, USA.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2227–2237, New Orleans, LA, USA.
- Kashyap Popat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2017. Where the truth lies: Explaining the credibility of emerging claims on the Web and social media. In *Proceedings of the Conference on World Wide Web, WWW '17*, pages 1003–1012, Perth, Australia.
- Benjamin Riedel, Isabelle Augenstein, Georgios P Spithourakis, and Sebastian Riedel. 2017. A simple but tough-to-beat baseline for the Fake News Challenge stance detection task. *ArXiv:1707.03264*.
- Parinaz Sobhani, Diana Inkpen, and Xiaodan Zhu. 2017. A dataset for multi-target stance detection. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2*, pages 551–557, Valencia, Spain.
- Karen Spärck Jones. 2004. Idf term weighting and ir research lessons. *Journal of documentation*, 60(5):521–523.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. End-to-end memory networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, pages 2440–2448, Montreal, Canada.
- Ming Tan, Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2016. Improved representation learning for question answer matching. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 464–473, Berlin, Germany.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT '18*, pages 809–819, New Orleans, LA, USA.
- Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. 2017. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR '17*, Honolulu, HI, USA.
- Andreas Vlachos and Sebastian Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proceedings of the ACL Workshop on Language Technologies and Computational Social Science*, pages 18–22, Baltimore, MD, USA.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.
- William Yang Wang. 2017. “liar, liar pants on fire”: A new benchmark dataset for fake news detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL '17*, pages 422–426, Vancouver, Canada.
- Caiming Xiong, Stephen Merity, and Richard Socher. 2016. Dynamic memory networks for visual and textual question answering. In *Proceedings of the 33rd International Conference on Machine Learning, ICML '16*, pages 2397–2406, New York, NY, USA.
- Brian Xu, Mitra Mohtarami, and James Glass. 2018. Adversarial domain adaptation for stance detection. In *Proceedings of the NeurIPS Workshop on Continual Learning*, Montreal, Canada.
- Guido Zarrella and Amy Marsh. 2016. MITRE at SemEval-2016 Task 6: Transfer learning for stance detection. In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval '16*, pages 458–463, San Diego, CA, USA.
- Yifan Zhang, Giovanni Da San Martino, Alberto Barrón-Cedeño, Salvatore Romeo, Jisun An, Hae-woon Kwak, Todor Staykovski, Israa Jaradat, Georgi Karadzhov, Ramy Baly, Kareem Darwish, James Glass, and Preslav Nakov. 2019. Tanbih: Get to know what you are reading. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (Demonstrations), EMNLP '19*, Hong Kong, China.