

Rumour Detection via Zero-shot Cross-lingual Transfer Learning

Lin Tian¹, Xiuzhen Zhang^{1*}, and Jey Han Lau²

¹ RMIT University, Australia

s3795533@student.rmit.edu.au, xiuzhen.zhang@rmit.edu.au

² The University of Melbourne, Australia

jeyhan.lau@gmail.com

Abstract. Most rumour detection models for social media are designed for one specific language (mostly English). There are over 40 languages on Twitter and most languages lack annotated resources to build rumour detection models. In this paper we propose a zero-shot cross-lingual transfer learning framework that can adapt a rumour detection model trained for a *source language* to another *target language*. Our framework utilises pretrained multilingual language models (e.g. multilingual BERT) and a self-training loop to iteratively bootstrap the creation of “silver labels” in the target language to adapt the model from the source language to the target language. We evaluate our methodology on English and Chinese rumour datasets and demonstrate that our model substantially outperforms competitive benchmarks in *both* source and target language rumour detection.

Keywords: Rumour Detection · Cross-lingual Transfer · Zero-shot.

1 Introduction

Online social media platforms provide an alternative means for the general public to access information. The ease of creating a social media account has the implication that rumours — *stories or statements with unverified truth value* [1] — can be fabricated by users and spread quickly on the platform.

To combat misinformation on social media, one may rely on fact checking websites such as `snopes.com` and `emergent.info` to dispel popular rumours. Although manual evaluation is the most reliable way of identifying rumours, it is time-consuming.

Automatic rumour detection is therefore desirable [14, 36]. Content-based methods focus on rumour detection using the textual content of messages and user comments. Feature-based models exploit features other than text content, such as author information and network propagation features, for rumour detection [19,

* Corresponding author

Table 1: An illustration of a COVID-19 rumour being circulated in English, French and Italian on Twitter.

Date	Language	Tweet
04-02-2020	English	Bill Gates admits the vaccine will no doubt kill 700000 people. The virus so far has killed circa 300000 globally. Can anyone explain to me why you would take a vaccine that kills more people than the virus it's designed to cure?
17-04-2020	French	et si bill gates etait le seul manipulateur de ce virus.. il veut moinsvde gens sur terre. veut vous vacciner et parle de pandemie depuis des années c est quand meme fou cette citation, non? #covid #BillGates
06-05-2020	Italian	Bill Gates:"la cosa più urgente nel mondo ora è il vaccino contro il Covid-19." I bambini africani che hanno ricevuto i vaccini di Bill Gates o sono morti o sono diventati epilettici. I vaccini di Bill Gates sono più pericolosi di qualsiasi coronavirus. #BillGates #Coronavirus

20, 22]. Most rumour detection models, however, are built for English [29, 30], and most annotated rumour datasets are also in English [13, 20, 28].

Rumours can spread in different languages and across languages. Table 1 shows an example (untruthful) rumour about Bill Gates circulated on Twitter during the COVID-19 pandemic.³ The rumour is found not only in English but also in French and Italian.⁴

There are over 40 languages on Twitter⁵ and most languages lack annotated resources for building rumour detection models. Although we have seen recent successes with deep learning based approaches for rumour detection [4, 15, 18, 19, 29, 39] most systems are monolingual and require annotated data to train a new model for a different language.

In this paper, we propose a zero-shot cross-lingual transfer learning framework for building a rumour detection system without requiring annotated data for a new language. Our system is cross-lingual in the sense that it can detect rumours in two languages based on one model. Our framework first fine-tunes a multilingual pretrained language model (e.g. multilingual BERT) for rumour detection using annotated data for a source language (e.g. English), and then uses it to classify rumours on another target language (zero-shot prediction) to create "silver" rumour labels for the target language. We then use these silver labels to fine-tune the multilingual model further to adapt it to the target language.

At its core, our framework is based on MultiFiT [8] which uses a multilingual model (LASER [2]) to perform zero-shot cross-lingual transfer from one language

³ <https://www.bbc.com/news/52847648>

⁴ The following article clarifies several rumours surrounding Bill Gates: <https://www.reuters.com/article/uk-factcheck-gates-idUSKBN2613CK>.

⁵ https://semicast.com/downloads/Semicast_Half_of_messages_on_Twitter_are_not_in_English.20100224.pdf

to another. An important difference is that we additionally introduce a self-training loop — which iteratively refines the quality of the silver labels — that can substantially improve rumour detection in the *target language*. Most interestingly, we also found that if we include the original gold labels in the source language in the self-training loop, detection performance in the *source language* can also be improved, creating a rumour detection system that excels in *both* source and target language detection.

To summarise, our contributions are: (1) we extend MultiFiT, a zero-shot cross-lingual transfer learning framework by introducing a self-training loop to build a cross-lingual model; and (2) we apply the proposed framework to the task of rumour detection, and found that our model substantially outperforms benchmark systems in both source and target language rumour detection.

2 Related Work

Rumour detection approaches can be divided into two major categories according to the types of data used: text-based and non-text based. Text-based methods focus on rumour detection using the textual content, which may include the original source document/message and user comments/replies. A study [18] proposed a recursive neural network model to detect rumours. Their model first clusters tweets by topics and then performs rumour detection at the topic level. Another study [29] introduced linguistic features to represent writing styles and other features based on sensational headlines from Twitter and to detect misinformation. To detect rumours as early as possible, a study [39] incorporated reinforcement learning to dynamically decide how many responses are needed to classify a rumour. Some other study [30] explored the relationship between a source tweet and its comments by transferring stance prediction model to classify the veracity of a rumour. Non-text-based methods utilise features such as user profiles or propagation patterns for rumour detection [15, 19]. In this paper, we adopt the text-based approach to rumour detection.

Most studies on rumour detection focus on a specific social media platform or language (typically English). Still there are a few exceptions that explore cross-domain/cross-lingual rumour detection or related tasks. A study [31] proposed a set of 10 hand-crafted cross-lingual and cross-platform features for rumour detection by capturing the similarity and agreement between online posts from different social media platforms. Another study [24] introduced a contrastive learning-based model for cross-lingual stance detection using memory networks. Different to these studies, we specifically focus on how to transfer learned knowledge from a source language to a target language for automatic rumour detection.

Transfer learning has been successfully applied to many natural language processing (NLP) tasks, where modern pretrained language models (e.g. BERT) are fine-tuned with annotated data for down-stream tasks [7, 16, 38]. Multilingual pretrained language models have also been explored. For example, BERT has a multilingual version trained using 104 languages of Wikipedia.⁶ A study [6]

⁶ <https://github.com/google-research/bert/blob/master/multilingual.md>.

incorporated RoBERTa’s training procedure to pretrain a multilingual language model that produces sentence embeddings for 100 languages. It is found that multilingual BERT is surprisingly good at zero-shot cross-lingual transfer [25]; in other words, it can be fine-tuned for a particular task in one language and used to make predictions in another language without any further training. MultiFiT [8] was recently proposed, where a zero-shot cross-lingual transfer framework uses predicted labels from a fine-tuned multilingual model to train a monolingual model on the same task in a target language; the transfer learning objective is only to optimise the model for the target language. Different from MultiFiT, our objective is to optimise models for both the target and source languages.

Self-training [27] is an early semi-supervised learning approach that has been explored for a variety of NLP tasks, such as neural machine translation [11], semantic segmentation [40]. Self-training involves teacher and student models, where the teacher model is trained with labelled data and then used to make predictions on unlabelled data to create more training data for training a student model. The process is repeated for several iterations with the student model replacing the original teacher model at the end of each iteration, and through iterative refinement of the predicted labels the student model improves over time. We apply self-training in a novel way to fine-tune pre-trained multilingual language models for cross-lingual rumour detection, and show that the student model improves over time during the transfer.

3 Methodology

We are interested in the task of rumour detection, and particularly how to do zero-shot cross-lingual transfer to build a cross-lingual rumour detection model. That is, we assume we have labelled rumours in one language (*source*) where we can build a supervised rumour detection model, and the goal is to transfer the model to detect rumours in a second language (*target*) without any labelled data in that second language. After transfer, it should have the ability to detect rumours in both languages (hence a multilingual model). We first describe the rumour classifier in Section 3.1, and return to detail the cross-lingual transfer learning framework in Section 3.2.

3.1 Rumour Classifier

We focus on binary rumour detection, and follow previous studies to classify whether a microblog post constitutes a rumour or not based on crowd comments [18, 30, 39].

Given an initial post s_i and its reactions r_i ,⁷ we feed them to a pretrained multilingual language model (we use multilingual BERT [7] and XML-RoBERTa

⁷ Reactions are replies and quotes. r_i represents all reactions that can fit the maximum sequence length (384) for the pretrained model, concatenated together as a long string.

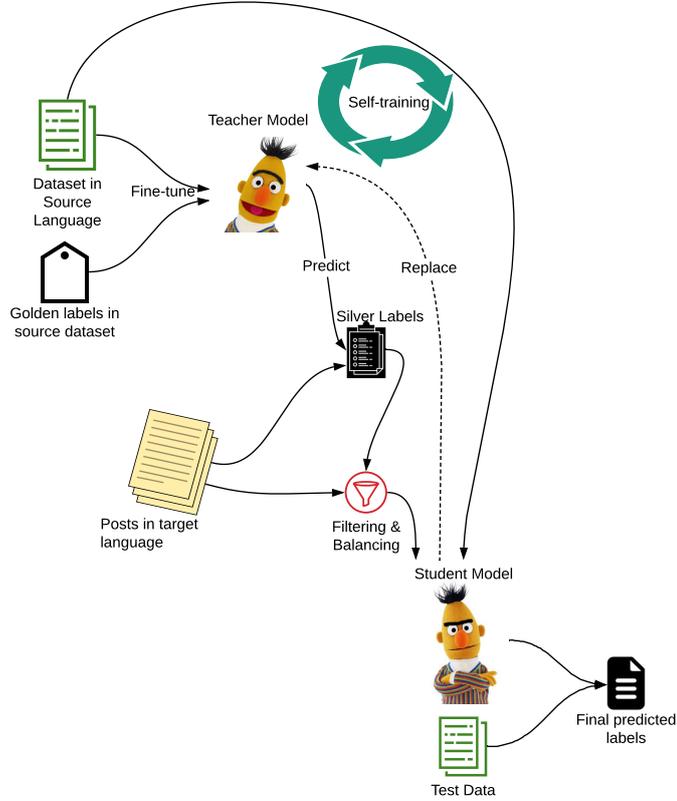


Fig. 1: Proposed cross-lingual transfer framework.

[6] in our experiments) as:⁸

$$[CLS] + s_i + [SEP] + r_i + [SEP]$$

where $[CLS]$ and $[SEP]$ are special symbols used for classification and separating sequences [7].

We then take the contextual embedding of $[CLS]$ ($h_{[CLS]}$) and feed it to a fully-connected layer to perform binary classification of the rumour.

$$y_i = \text{softmax}(W_i h_{[CLS]} + b_s)$$

Given ground truth rumour labels, the model is fine-tuned with standard binary cross-entropy loss. All parameters are updated except for the word embeddings (rationale detailed in the following section).

⁸ For XLM-RoBERTa, we have 2 $[SEP]$ symbols between s_i and r_i , following https://huggingface.co/transformers/model_doc/xlmroberta.html#transformers.XLMRobertaTokenizer.build_inputs_with_special_tokens.

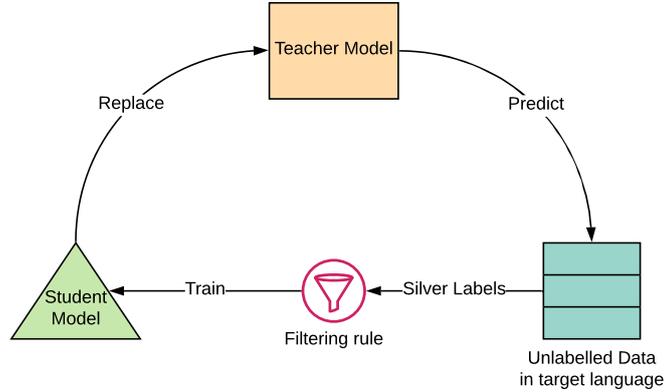


Fig. 2: Self-training loop.

3.2 Cross-lingual Transfer

Our zero-shot cross-lingual transfer learning framework is based on MultiFiT [8]. MultiFiT works by first fine-tuning a multilingual model (e.g. LASER [2] is used in the original paper) for a task in a source language, and then applying it (zero-shot) to the same task in a target language to create silver labels. These silver labels are then used to fine-tune a monolingual model in the target language. MultiFiT is shown to substantially improve document classification compared to zero-shot predictions by a series of multilingual models trained using only gold labels in the source language.⁹

We present our zero-shot cross-lingual transfer learning framework in Figure 1. One key addition that we make is a self-training loop that *iteratively refines the quality of the adapted model*. In the original MultiFiT framework, the teacher model is a multilingual model, and the student model is a monolingual model in the target language. As we are interested in multilingual rumour detection, the student model is a multilingual model in our case, although in our experiments (Section 4.3) we also present variations where the student model is a monolingual model.

Figure 2 illustrates the self-training loop. The student model is initialised using the teacher model (so both are multilingual models). Once the student is trained, the teacher model in the next iteration will be replaced by the student model.

To reduce noise in the silver labels, we introduce a filtering and balancing procedure in the self-training loop. The procedure was originally introduced to image classification and shown to improve performance [34]. With the filtering

⁹ Silver labels refer to the predicted labels in the target language, while gold labels refer to the real labels in the source language.

procedure, instances with prediction confidence/probability lower than a threshold p are filtered. The balancing procedure effectively drops some high confidence instances that pass the threshold to ensure that an equal number of positive (rumour) and negative (non-rumour) instances.

Following [10], we perform adaptive pretraining on the teacher model before fine-tuning it for the rumour detection task. That is, we take the off-the-shelf pretrained multilingual model and further pretrain it using the masked language model objective on data in our rumour detection/social media domain. In terms of pretraining data we use both the unlabelled rumour detection data (“task adaptive”) and externally crawled microblog posts (“domain adaptive”) in the target language.

The degree of overlap in terms of vocabulary between the source and target language varies depending on the language pair. If the overlap is low, after fine-tuning the source language subword embeddings would have shifted while the target language subword embeddings remained the same (due to no updates), creating a synchronisation problem between the subword embeddings and intermediate layers. We solve this issue by freezing the subword embeddings when we first fine-tune the teacher model; subsequent fine-tuning in the self-training loop, however, updates all parameters. We present ablation tests to demonstrate the importance of doing this in Section 4.4.

Aiming for a model that performs well for both the target and source languages, we also introduce gold labels in the source language during self-training, i.e. we train the student using both the silver labels in the target language and the gold labels in the source language. This approach produces a well-balanced rumour detection model that performs well in both source and target languages, as we will see in Section 4.3.

4 Experiments and Results

We evaluated our cross-lingual transfer learning framework for rumour detection using three English and Chinese datasets. We formulate the problem as a binary classification task to distinguish rumours from non-rumours.

4.1 Datasets

Two English datasets Twitter15/16 [20] and PHEME [13], and one Chinese dataset WEIBO [18] were used in our experiments. PHEME and WEIBO have two class labels, rumour and non-rumour. For the Chinese WEIBO dataset, rumours are defined as “a set of known rumours from the Sina community management center (<http://service.account.weibo.com>), which reports various misinformation” [18]. The original Twitter15/16 dataset [20] has four classes, true rumour, false rumour, unverified rumour and non-rumour. We therefore extract tweets with labels “false rumour” and “non-rumour” from Twitter15/16 to match

Table 2: Rumour datasets.

	T15/16	PHEME	WEIBO
#initial posts	1,154	2,246	4,664
#all posts	182,535	29,387	3,805,656
#users	122,437	20,529	2,746,818
#rumours	575	1,123	2,313
#non-rumours	579	1,123	2,351
Avg. # of reactions	279	26	247
Max. # of reactions	3,145	289	2,313
Min. # of reactions	74	12	10

the definition of rumours and non-rumours of WEIBO and use the extracted data for experiments. Table 2 shows statistics of the experiment datasets.

To ensure fair comparison of the performance across all models, for each dataset we reserved 20% data as test and we split the rest in a ratio of 4:1 for training and validation partitions. The validation set was used for hyperparameter tuning and early-stopping. For the PHEME dataset, to be consistent with the experiment set up in the literature, we followed the 5-fold split from [21]. For adaptive pretraining (Section 3.2), we used an external set of microblogs data for English (1.6M posts; [28]) and Chinese (39K posts).¹⁰

4.2 Experiment Setup

We used multilingual BERT [7] and XLM-RoBERTa [6] for the multilingual models, and implemented in PyTorch using the HuggingFace Libraries.¹¹

For adaptive pretraining, we set batch size=8. For the fine-tuning, we set batch size=16, maximum token length=384, and dropout rate=0.1. Training epochs vary between 3–5 and learning rate in the range of $\{1e-5, 2e-5, 5e-5\}$; the best configuration is chosen based on the development data. We also tuned the number of self-training iterations and p , the threshold for filtering silver labels (Section 3.2), based on development.¹² All experiments were conducted using $1 \times V100$ GPU.

4.3 Results

For our results, we show cross-lingual transfer performance from English to Chinese and vice versa. As we have two English datasets (T15/16 and PHEME) and one Chinese dataset (WEIBO), we have four sets of results in total: T15/16→WEIBO,

¹⁰ <https://archive.ics.uci.edu/ml/datasets/microblogPCU>

¹¹ <https://github.com/huggingface>.

¹² For p we search in the range of 0.94–0.96.

Table 3: Rumour detection results (Accuracy (%)) for English to Chinese transfer. Each result is an average over 3 runs, and subscript denotes standard deviation. monoBERT is a Chinese BERT model in this case. Boldfont indicates optimal zero-shot performance.

Model	T15/16→WEIBO		PHEME→WEIBO	
	Source	Target	Source	Target
Supervised				
multiBERT+source	95.8 _{0.1}	—	83.7 _{0.5}	—
multiBERT+target	—	93.9 _{0.2}	—	93.9 _{0.2}
multiBERT+both	94.8 _{0.1}	93.0 _{0.3}	82.1 _{0.8}	95.2 _{0.2}
XLMR+source	96.3 _{0.4}	—	82.8 _{0.5}	—
XLMR+target	—	94.8 _{0.1}	—	94.8 _{0.1}
XLMR+both	95.5 _{0.1}	92.2 _{0.2}	85.8 _{1.9}	95.4 _{0.1}
[18]+source	83.5 _{0.7}	—	80.8 _{0.4}	—
[15]+source	85.4 _{0.4}	—	64.5 _{1.0}	—
[30]+source	87.2 _{0.9}	—	86.7 _{1.5}	—
[4]+source	96.3 _{0.7}	—	—	—
Zero-shot				
multiBERT	—	64.3 _{2.1}	—	65.9 _{1.1}
XLMR	—	64.7 _{1.1}	—	68.1 _{1.0}
MF [8]	—	70.6 _{0.4}	—	61.1 _{1.0}
MF-monoBERT	—	67.5 _{0.5}	—	68.2 _{0.4}
MF-monoBERT+ST	—	81.3 _{0.1}	—	79.0 _{0.2}
MF-multiBERT+ST	61.3 _{0.4}	78.6 _{3.9}	66.3 _{1.5}	72.6 _{1.9}
MF-multiBERT+ST+GL	96.6 _{0.2}	78.3 _{0.8}	83.0 _{0.5}	74.3 _{3.3}
MF-XLMR+ST	57.6 _{1.0}	81.2 _{0.1}	62.1 _{1.3}	77.4 _{0.3}
MF-XLMR+ST+GL	96.2 _{0.1}	80.2 _{0.2}	85.3 _{0.7}	77.2 _{0.8}

PHEME→WEIBO, WEIBO→T15/16 and WEIBO→PHEME. We evaluate rumour detection performance using accuracy, and present the English→Chinese results in Table 3 and Chinese→English in Table 4 respectively. All performance is an average over 3 runs with different random seeds.

We include both supervised and zero-shot baselines in our experiments. For the supervised benchmarks, we trained multilingual BERT and XLM-RoBERTa using: (1) source labels; (2) target labels; and (3) both source and target labels. The next set of supervised models are state-of-the-art monolingual rumour detection models: (1) [18] is a neural model that processes the initial post and crowd comments with a 2-layer gated recurrent units; (2) [15] uses recurrent and convolutional networks to model user metadata (e.g. followers count) in the

Table 4: Rumour detection results (Accuracy (%)) for Chinese to English transfer. Each result is an average over 3 runs, and subscript denotes standard deviation. monoBERT is an English BERT model here.

Model	WEIBO→T15/16		WEIBO→PHEME	
	Source	Target	Source	Target
Supervised				
multiBERT+source	93.9 _{0.1}	—	93.9 _{0.2}	—
multiBERT+target	—	95.8 _{0.1}	—	83.7 _{0.5}
multiBERT+both	93.0 _{0.3}	94.8 _{0.1}	95.2 _{0.2}	82.1 _{0.8}
XLMR+source	94.8 _{0.1}	—	94.8 _{0.1}	—
XLMR+target	—	96.3 _{0.4}	—	82.8 _{0.5}
XLMR+both	92.2 _{0.2}	95.5 _{0.1}	95.4 _{0.1}	85.8 _{1.9}
[18]+source	91.0 _{0.1}	—	91.0 _{0.1}	—
[15]+source	92.1 _{0.2}	—	92.1 _{0.2}	—
[4]+source	96.1 _{0.4}	—	96.1 _{0.4}	—
Zero-shot				
multiBERT	—	60.8 _{1.3}	—	67.2 _{0.4}
XLMR	—	73.9 _{0.8}	—	69.0 _{1.5}
MF [8]	—	73.4 _{1.4}	—	64.1 _{2.0}
MF-monoBERT	—	64.7 _{0.1}	—	70.7 _{0.7}
MF-monoBERT+ST	—	85.7 _{0.4}	—	78.9 _{0.6}
MF-multiBERT+ST	55.1 _{1.8}	82.2 _{1.2}	66.0 _{1.0}	72.6 _{1.5}
MF-multiBERT+ST+GL	97.0 _{0.1}	80.9 _{1.5}	95.8 _{0.5}	73.4 _{0.4}
MF-XLMR+ST	52.4 _{0.3}	83.0 _{1.0}	62.7 _{0.5}	75.4 _{0.4}
MF-XLMR+ST+GL	97.6 _{0.1}	81.3 _{0.1}	95.9 _{0.5}	77.9 _{1.1}

crowd responses;¹³ (3) [30] uses BERT to encode comments (like our model) but it is pre-trained with stance annotations; and (4) [4] uses bidirectional graph convolutional networks to model crowd responses in the propagation path. Note that we only have English results (T15/16 and PHEME) for [30] as it uses stance annotations from SemEval-2016 [23], and only T15/16 and WEIBO results for [4] as PHEME does not have the propagation network structure. To get user metadata for [15], we crawled user profiles via the Twitter API.¹⁴

For the zero-shot baselines, multilingual BERT and XLM-RoBERTa were trained using the source labels and applied to the target language (zero-shot predictions); subword embeddings are frozen during fine-tuning for these zero-shot models. We also include the original MultiFiT model [8], which uses LASER [2] as

¹³ Following the original paper, only a maximum of 100 users are included.

¹⁴ <https://developer.twitter.com/en/docs/twitter-api/v1>

Table 5: Rumour detection results (F_1 score (%)) for both the source and target languages. “R” and “NR” denote the rumour and non-rumour classes respectively.

	T15/16→WEIBO		PHEME→WEIBO		WEIBO→T15/16		WEIBO→PHEME	
	Source	Target	Source	Target	Source	Target	Source	Target
MF-multiBERT+ST+GL								
R	96.6	79.1	84.2	75.5	96.9	83.3	94.3	75.2
NR	96.1	77.2	79.5	73.1	97.0	76.8	94.4	70.3
MF-XLMR+ST+GL								
R	96.6	83.1	85.2	81.6	97.4	82.9	96.9	74.4
NR	95.6	76.1	85.7	74.5	96.8	81.1	95.0	81.4

the multilingual model (teacher) and a pretrained quasi-recurrent neural network language model [5] as the monolingual model (student).¹⁵

We first look at the supervised results. XLM-RoBERTa (“XLMR”) is generally better (marginally) than multilingual BERT (“multiBERT”). In comparison, for the monolingual rumour detection models, [4] has the best performance overall (which uses network structure in addition to crowd comments), although XLM-RoBERTa and multilingual BERT are not far behind.

Next we look at the zero-shot results. Here we first focus on target performance and baseline models. The zero-shot models (“multiBERT” and XLMR”) outperform the MultiFiT baseline (“MF”) in 2–3 out of 4 cases, challenging the original findings in [8]. When we replace the teacher model with multilingual BERT and the student model with monolingual BERT (“MF-monoBERT”), we found mixed results compared to MultiFiT (“MF”): 2 cases improve but the other 2 worsen. When we incorporate the self-training loop (“MF-monoBERT+ST”), however, we see marked improvement in all cases — the largest improvement is seen in WEIBO→T15/16 (Chinese to English, Table 4), from 64.7% to 85.7% — demonstrating the benefits of iteratively refining the transferred model. These results set a new state-of-the-art for zero-shot cross-lingual transfer learning for our English and Chinese rumour detection datasets. That said, there is still a significant gap (10+ accuracy points) compared to supervised models, but as we see in Section 4.6 the gap diminishes quickly as we introduce some ground truth labels in the target domain.

We now discuss the results when we use a multilingual model for the student model, i.e. replacing it with either multilingual BERT (“MF-multiBERT+ST”) or XLM-RoBERTa (“MF-XLMR+ST”), which turns it into a *multilingual* rumour detection system (i.e. after fine-tuned it can detect rumours in both source and target language). Similar to the supervised results, we see that the latter (“MF-XLMR+ST”) is a generally better multilingual model. Comparing our best multilingual student model (“MF-XLMR+ST”) to the monolingual student model

¹⁵ The monolingual student model is pretrained using Wikipedia in the target language.

(“MF-monoBERT+ST”) we see only a small drop in the target performance (about 1–4 accuracy points depending on domain), demonstrating that the multilingual rumour detection system is competitive to the monolingual detection system in the target language.

For the source performance, we see a substantial drop (20–40 accuracy points) after cross-lingual transfer (e.g. “XLMR+source” vs. “MF-XLMR+ST”), implying there is catastrophic forgetting [9, 12, 32, 35] — the phenomenon where adapted neural models “forget” and perform poorly in the original domain/task. When we incorporate gold labels in the source domain in the self-training loop (“MF-multiBERT+ST+GL” or “MF-XLMR+ST+GL”), we found a surprising observation: not only was catastrophic forgetting overcome, but the source performance actually surpasses some supervised monolingual models, e.g. “MF-multiBERT+ST+GL” and “MF-XLMR+ST+GL” outperform [4] in T15/16 (96.6% vs. 96.3%) and WEIBO (97.6% vs. 96.1%) respectively, creating a new state-of-the-art for rumour detection in these two domains. One explanation is that the transfer learning framework maybe functioning like a unique data augmentation technique that creates additional data in a different language (unique in the sense it works only for improving multilingual models). Note that incorporating the gold labels generally does not hurt the *target performance* — e.g. comparing “MF-XLMR+ST” with “MF-XLMR+ST+GL” we see a marginal dip in 2 cases, but in 2 other cases we see similar or improved performance — which shows that this is an effective approach for building multilingual models.

We further examine class-specific performance of our best models. The F_1 scores of rumour and non-rumour classes are presented in Table 5. For this binary classification task with relatively balanced class distributions, not surprisingly we observe that our models have reasonably good performance in both the rumour and non-rumour classes; lowest F_1 score is 70.3% of MF-multiBERT+ST+GL (Chinese to English transfer) for non-rumours in PHEME. That said, performance of the rumour class is generally better than that of the non-rumour class in both the source and target languages (the only exception is Chinese to English transfer on PHEME).

4.4 Adaptive Pretraining and Layer Freezing

To understand the impact of adaptive pretraining and layer freezing, we display zero-shot multilingual BERT results (test set) in Table 6. We can see that there are clear benefits for adaptive pretraining (top-3 vs. bottom-3 rows). For layer freezing, we have 3 options: no freezing (“ \emptyset ”), only freezing the subword embeddings (“*”) and freezing the first 3 layers (“***”). The second option (subword embedding frozen) consistently produces the best results (irrespective of whether adaptive pretraining is used), showing that this approach is effective in tackling the synchronisation issue (Section 3.2) that arises when we fine-tune a multilingual model on one language.

Table 6: Influence of adaptive pretraining (“Ad. Pt.”) and layer freezing (“Frz.”) for results (Accuracy (%)). “ \emptyset ” denotes no freezing of any layers; “*” freezing the subword embedding layer; and “**” freezing the first 3 layers.

Ad. Pt.	Frz.	T15/16 \rightarrow PHEME \rightarrow WEIBO \rightarrow WEIBO \rightarrow			
		WEIBO	WEIBO	T15/T16	PHEME
N	\emptyset	53.2	58.4	50.1	54.5
	*	61.3	60.0	52.6	63.8
	**	57.5	60.9	52.9	58.5
Y	\emptyset	56.6	61.8	50.9	61.6
	*	64.3	65.9	60.8	67.2
	**	60.3	63.8	55.0	61.8

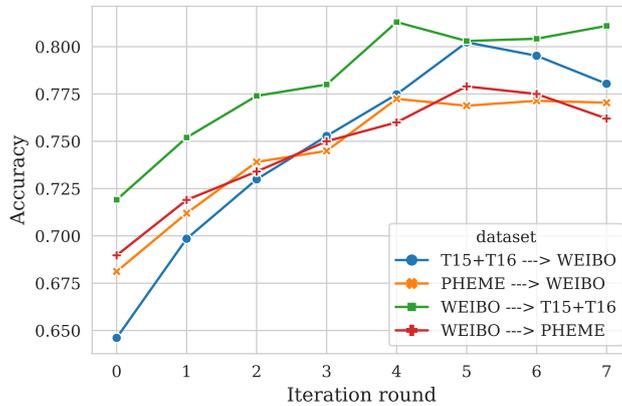


Fig. 3: Accuracy over iteration during self-training.

4.5 Self-training

To measure the influence of the self-training loop, we present target performance (test set) of our multilingual model (“MF-XLMR+ST+GL”) over different iterations in the self-training loop in Figure 3. We can see the performance improves rapidly in the first few iterations, and gradually converges after 4–7 iterations. These results reveal the importance of refining the model over multiple iterations during cross-lingual transfer.

4.6 Semi-supervised Learning

Here we explore feeding a proportion of ground truth labels in the target domain to our zero-shot model (“MF-XLMR+ST+GL”) and compare it to supervised

Table 7: T15/16→WEIBO results (Accuracy(%)) as we incorporate more ground truth target labels (“GT Label”).

% GT Label	Supervised	Zero-shot
0%	—	80.2
20%	79.8	86.3
40%	83.3	89.2
60%	89.3	92.9
80%	91.0	93.5
100%	92.2	—

multilingual model (“XLMR+both”). We present T15/16→WEIBO results (test set) in Table 7. We can see that the gap shrinks by more than half (12.0 to 5.9 accuracy difference) with just 20% ground truth target label. In general our unsupervised cross-lingual approach is also about 20% more data efficient (e.g. supervised accuracy@40% \approx unsupervised accuracy@20%). Interestingly, with 60% ground truth our model outperforms the fully supervised model.

5 Discussion and Conclusions

One criticism of the iterative self-training loop is that it suffers from poor initial prediction which could lead to a vicious cycle that further degrades the student model. The poor initial predictions concern appears to less of a problem in our task, as the pure zero-shot models (i.e. without self-training) appear to do reasonably well when transferred to a new language, indicating that the pretrained multilingual models (e.g. XLMR) are sufficiently robust. By further injecting the gold labels from the source domain during self-training, we hypothesise that it could also serve as a form of regularisation to prevent continuous degradation if the initial predictions were poor. Also, although our proposed transfer learning framework has only been applied to multilingual rumour detection, the architecture of the framework is general and applicable to other tasks.

To conclude, we proposed a zero-shot cross-lingual transfer learning framework to build a multilingual rumour detection model using only labels from one language. Our framework introduces: (1) a novel self-training loop that iteratively refines the multilingual model; and (2) ground truth labels in the source language during cross-lingual transfer. Our zero-shot multilingual model produces strong rumour detection performance in both source and target language.

Acknowledgments

This research is supported in part by the Australian Research Council Discovery Project DP200101441.

References

1. Allport, G.W., Postman, L.: The psychology of rumor. (1947)
2. Artetxe, M., Schwenk, H.: Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. Transactions of the Association for Computational Linguistics. (2019)
3. Bengio, Y.: Deep learning of representations for unsupervised and transfer learning. In: Proceedings of ICML workshop on unsupervised and transfer learning. (2012)
4. Bian, T., Xiao, X., Xu, T., Zhao, P., Huang, W., Rong, Y., Huang, J.: Rumor detection on social media with bi-directional graph convolutional networks. In: AAAI (2020)
5. Bradbury, J., Merity, S., Xiong, C., Socher, R.: Quasi-recurrent neural networks. (2016)
6. Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., Stoyanov, V.: Unsupervised cross-lingual representation learning at scale. (2019)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: NAACL (2019)
8. Eisenschlos, J., Ruder, S., Czapla, P., Kardas, M., Gugger, S., Howard, J.: MultiFiT: Efficient multi-lingual language model fine-tuning. (2019)
9. French, R.M.: Catastrophic forgetting in connectionist networks. Trends in cognitive sciences. (1999)
10. Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., Smith, N.A.: Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. In: ACL (2020)
11. He, J., Gu, J., Shen, J., Ranzato, M.: Revisiting self-training for neural sequence generation. (2019)
12. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., *et. al*: Overcoming catastrophic forgetting in neural networks. Proceedings of the national academy of sciences. (2017)
13. Kochkina, E., Liakata, M., Zubiaga, A.: All-in-one: Multi-task Learning for Rumour Verification. In: COLING (2018)
14. Liu, X., Nourbakhsh, A., Li, Q., Fang, R., Shah, S.: Real-time rumor debunking on Twitter. In: CIKM (2015)
15. Liu, Y., Wu, Y.F.B.: Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In: AAAI (2018)
16. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. (2019)
17. Long, Y., Lu, Q., Xiang, R., Li, M., Huang, C.R.: Fake news detection through multi-perspective speaker profiles. In: IJCNLP (2017)

18. Ma, J., Gao, W., Mitra, P., Kwon, S., Jansen, B.J., Wong, K.F., Cha, M.: Detecting rumors from microblogs with recurrent neural networks. In: IJCAI (2016)
19. Ma, J., Gao, W., Wei, Z., Lu, Y., Wong, K.F.: Detect rumors using time series of social context information on microblogging websites. In: CIKM (2015)
20. Ma, J., Gao, W., Wong, K.F.: Detect rumors in microblog posts using propagation structure via kernel learning. In: ACL (2017)
21. Ma, J., Gao, W., Wong, K.F.: Detect rumors on Twitter by promoting information campaigns with generative adversarial learning. In: WWW (2019)
22. Mendoza, M., Poblete, B., Castillo, C.: Twitter under crisis: Can we trust what we rt? In: Proceedings of the first workshop on social media analytics (2010)
23. Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., Cherry, C.: Semeval-2016 task 6: Detecting stance in tweets. In: Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016) (2016)
24. Mohtarami, M., Glass, J., Nakov, P.: Contrastive language adaptation for cross-lingual stance detection. In: EMNLP-IJCNLP (2019)
25. Pires, T., Schlinger, E., Garrette, D.: How multilingual is multilingual bert? (2019)
26. Ren, Y., Zhang, J.: HGAT: Hierarchical Graph Attention Network for Fake News Detection. (2020)
27. Scudder, H.: Probability of error of some adaptive pattern-recognition machines. IEEE Transactions on Information Theory. (1965)
28. Shu, K., Mahudeswaran, D., Wang, S., Lee, D., Liu, H.: Fakenewsnet: A data repository with news content, social context and dynamic information for studying fake news on social media. Big Data (2020)
29. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: A data mining perspective. ACM SIGKDD Explorations Newsletter (2017)
30. Tian, L., Zhang, X., Wang, Y., Liu, H.: Early detection of rumours on Twitter via stance transfer learning. In: ECIR (2020)
31. Wen, W., Su, S., Yu, Z.: Cross-lingual cross-platform rumor verification pivoting on multimedia content. In: EMNLP (2018)
32. Wiese, G., Weissenborn, D., Neves, M.: Neural domain adaptation for biomedical question answering. (2017)
33. Wu, L., Liu, H.: Tracing fake-news footprints: Characterizing social media messages by how they propagate. In: WSDM (2018)
34. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: CVPR (2020)
35. Xu, Y., Zhong, X., Yepes, A.J.J., Lau, J.H.: Forget me not: Reducing catastrophic forgetting for domain adaptation in reading comprehension. In: IJCNN (2020)
36. Yang, F., Liu, Y., Yu, X., Yang, M.: Automatic detection of rumor on sina weibo. In: Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics. (2012)
37. Yang, S., Shu, K., Wang, S., Gu, R., Wu, F., Liu, H.: Unsupervised fake news detection on social media: A generative approach. In: AAI (2019)
38. Yin, D., Meng, T., Chang, K.W.: Sentibert: A Transferable Transformer-Based Architecture for Compositional sentiment semantics. In: ACL (2020)
39. Zhou, K., Shu, C., Li, B., Lau, J.H.: Early rumour detection. In: NAACL (2019)
40. Zou, Y., Yu, Z., Vijaya Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: ECCV (2018)