

CONCRETE: Improving Cross-lingual Fact-checking with Cross-lingual Retrieval

Kung-Hsiang Huang ChengXiang Zhai Heng Ji

Department of Computer Science, University of Illinois Urbana-Champaign

{khhuang3, czhai, hengji}@illinois.edu

Abstract

Fact-checking has gained increasing attention due to the widespread of falsified information. Most fact-checking approaches focus on claims made in English only due to the data scarcity issue in other languages. The lack of fact-checking datasets in low-resource languages calls for an effective cross-lingual transfer technique for fact-checking. Additionally, trustworthy information in different languages can be complementary and helpful in verifying facts. To this end, we present the first fact-checking framework augmented with cross-lingual retrieval that aggregates evidence retrieved from multiple languages through a cross-lingual retriever. Given the absence of cross-lingual information retrieval datasets with claim-like queries, we train the retriever with our proposed *Cross-lingual Inverse Cloze Task (X-ICT)*, a self-supervised algorithm that creates training instances by translating the title of a passage. The goal for X-ICT is to learn cross-lingual retrieval in which the model learns to identify the passage corresponding to a given translated title. On the X-FACT dataset, our approach achieves 2.23% absolute F1 improvement in the zero-shot cross-lingual setup over prior systems. The source code and data are publicly available at <https://github.com/khuangaf/CONCRETE>.

1 Introduction

Fact-checking is an important task that assesses the veracity of a claim. This task has gained increasing attention due to the widespread mis- and disinformation that has a significant socioeconomic impact on our society (Scheufele and Krause, 2019; Pate et al., 2019; Fung et al., 2021; Guo et al., 2022; Wu et al., 2022; Fung et al., 2022; Huang et al., 2022). While fact-checking is mainly conducted manually, especially in the journalism industry, with more than hundred of millions of social media posts and millions of blog posts published per day (Hoang and Mothe, 2018; Djuraskovic,

2022), manual fact-checking is no longer feasible. Hence, we are urgently in need of reliable automated fact-checking approaches.

Most existing work develops fact-checking approaches on English-only corpus (Wang, 2017; Thorne et al., 2018; Augenstein et al., 2019; Wadden et al., 2020). One reason for this is the scarcity of fact-checking websites in other languages for constructing large enough non-English fact-checking datasets. Therefore, it is even more challenging to build fact-checkers for low-resource languages. A solution is to leverage high-resource languages with zero-shot cross-lingual transfer, where the model is trained on source languages of richer resources and directly tested on target languages of lower resources. Note that the sets of languages in the training set and the test sets are disjoint. With this technique, ground-truth labels for claims in low-resource languages are no longer required.

Few studies in the fact-checking literature have explored the cross-lingual setup. One line of work attempts to match an input claim with claims in other languages that have been verified (Kazemi et al., 2021, 2022). However, this approach fails when the claims have not been fact-checked in any language. Another line of work builds classifiers with multilingual language models. For example, Gupta and Srikumar (2021) utilize Google Search to obtain snippets that are relevant to the input claim and train a model based on mBERT (Devlin et al., 2019) in a cross-lingual setting. Although Google Search excels at retrieving relevant information from a given claim, it disregards the trustworthiness of the information being retrieved. When a claim is erroneous, the search results often contradict each other, which likely leads to incorrect predictions, as shown later in Section 5.3.

Motivated by these challenges, we propose CONCRETE, a Claim-oriented Cross-lingual Retriever that retrieves evidence from a trustworthy multi-

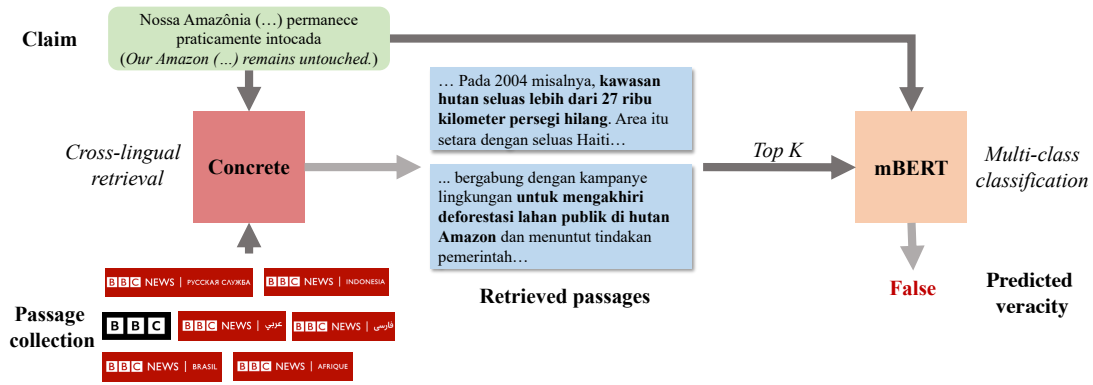


Figure 1: An overview of the proposed framework. Given a claim in arbitrary language, a cross-lingual retriever, CONCRETE retrieves relevant passages in *any languages*. The top- k relevant passages and the claim are then passed to our multilingual reader, mBERT, to predict the veracity of the claim.

lingual passage collection for fact-checking. This approach can handle region-specific claims as long as relevant evidence is presented in the passage collection, which is much more accessible compared to similar claims in other languages. In addition, since it does not rely on a black-box retrieval system, our approach provides the flexibility to include only trustworthy information in the passage collection. One major challenge for training such a retriever is the lack of multilingual information retrieval (IR) dataset with claim-like queries. To this end, we propose a self-supervised cross-lingual learning algorithm, Cross-lingual Inverse Cloze Task (X-ICT), to learn the retriever based on pseudo-feedback. To mimic cross-lingual retrieval, we construct a pseudo-query for a given passage by translating its title into a randomly selected language. The objective for X-ICT is to identify the passage corresponding to a given translated title among all candidate passages. Since the title of news articles can often be regarded as a claim, this approach mitigates the domain discrepancy issue. As shown in Figure 1, our framework first performs cross-lingual retrieval to obtain evidence relevant to the input claim. Then, a multilingual reader takes in the retrieved evidence and the input claim to classify the veracity of the claim.

Our contributions can be summarized as follows:

- To the best of our knowledge, we present the first fact-checking framework augmented with cross-lingual retrieval that achieves state-of-the-art cross-lingual transfer performance on the X-FACT fact-checking task.
- We propose CONCRETE, a cross-lingual retriever with a bi-encoder architecture learned through a proposed self-supervised learning algorithm.

- Our experiments reveal that the distance between the input claim and the retrieved passages is strongly correlated with the performance.
- We collected a multilingual passage collection composed of reliable news articles in seven different languages. We have demonstrated that this corpus is effective for retrieval-augmented fact-checking.

2 Task Definitions

The input is a claim c in an arbitrary language and the corresponding metadata, such as claimer and claim date. Based on c , the retriever component of our model retrieves relevant passages p from a multilingual passage collection P where p can be in any language. Then, the reader component takes in the claim c , the corresponding metadata, and relevant passages p to predict the veracity of c . Note that the use of passage collection P to aid in fact-checking is a modeling choice, as no grounded evidence is available.

We aim to build a passage collection that does not contain erroneous information. Hence, we construct P by crawling 49,000 articles published in 7 languages¹ between September 2016 and December 2022 from a trustworthy news media, bbc.com, and split each article into passages. Each passage contains at most 100 tokens, following Karpukhin et al. (2020). This results in a total of 347,557 passages. The collected passage collection has been made publicly available in the link mentioned in the Abstract.

¹We consider Arabic, Russian, Indonesian, Persian, French, and Portuguese.

3 Proposed Method

Our framework is a pipeline consisting of two components: (1) CONCRETE, a claim-oriented cross-lingual retriever that retrieves relevant passages from a multilingual passage collection, and (2) a multilingual reader that determines the veracity of a claim based on the compatibility of the claim and the passages retrieved. Figure 1 shows an overview of the proposed method. The following sections describe each component in detail.

3.1 CONCRETE

CONCRETE takes in a claim c in an *arbitrary language* as input and retrieves k relevant passages $p = \{p_0, \dots, p_k\}$ from a multilingual passage collection \mathcal{P} , where the retrieved passages can be in *any language*. We adapt mDPR (Asai et al., 2021b), a multilingual retriever, with the proposed self-supervised learning algorithm, Cross-lingual Inverse Cloze Task, for claim-like queries.

mDPR mDPR is a multilingual version of the Dense Passage Retriever (DPR) (Karpukhin et al., 2020). Similar to DPR, mDPR is a bi-encoder architecture that computes the relevance score between a query and a passage with the inner product of the corresponding representations. In addition to being trained on existing multilingual question answering datasets, mDPR also learns from additional samples mined from Wikipedia and labeled with the proposed answer generator discussed in Asai et al. (2021b). However, since mDPR was trained on datasets where queries are questions instead of claims, domain mismatch becomes an issue if we directly apply mDPR to our task.

Cross-lingual Inverse Cloze Task To address the domain discrepancy problem discussed in the previous paragraph, a naive solution is to fine-tune mDPR on multilingual IR datasets with claim-like queries. Unfortunately, such datasets are not available. Motivated by the Inverse Cloze Task (Lee et al., 2019), which was used to warm start the retriever by tasking a retriever to predict the context given a randomly sampled sentence, we propose Cross-lingual Inverse Cloze Task (X-ICT) by extending ICT to a cross-lingual setup. Specifically, we made two major modifications to the original ICT.

First, instead of randomly sampling a sentence as the pseudo query, we treat the title of each passage as the pseudo query as these titles are often

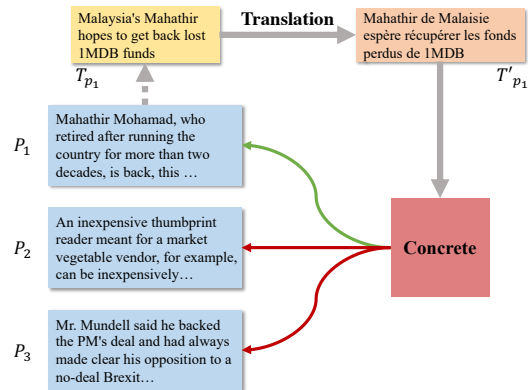


Figure 2: An illustration of X-ICT. Given a passage p_1 , we find its title T_{p_1} and translates it into a different language T'_{p_1} to mimic cross-lingual retrieval. The goal for the cross-lingual retriever, CONCRETE, is to select the correct passage p_1 based on the translated title T'_{p_1} among all the passages in the same batch.

claim-like sentences. The positive passages for a given title are the passages derived from the same article. The biggest advantage of this approach is that there is very little domain mismatch between using claims as queries and using titles as queries. Therefore, the retriever optimized with our proposed X-ICT can be directly applied to claim-oriented downstream tasks without further fine-tuning, indicating that the optimization for the downstream tasks can be more efficient. Second, to mimic cross-lingual retrieval, we use mBART-50 (Tang et al., 2020), a machine translation model, to translate the title into a target language.

Formally, in CONCRETE, two dense encoders, $E_C(\cdot)$ and $E_P(\cdot)$, are used to represent claims and passages as d -dimensional vectors. The similarity between a claim and a passage is defined as the dot product between their dense vectors,

$$\text{sim}(c, p_i) = E_C(c)^\top E_P(p_i). \quad (1)$$

In X-ICT, we form positive claim-passage pairs by constructing the translation of the corresponding title T_{p_i} of a given passage p_i . i.e. (T'_{p_i}, p_i) where T'_{p_i} is the original title T_{p_i} translated into a different language using mBART-50. We treat other passages in the same batch as negative samples. The retriever is trained by optimizing the negative log likelihood,

$$P(p_i | T'_{p_i}) = \frac{\exp(\text{sim}(T'_{p_i}, p_i))}{\sum_{p_j \in \text{BATCH}} \exp(\text{sim}(T'_{p_i}, p_j))} \quad (2)$$

$$\mathcal{L}_{\text{X-ICT}} = -\log \sum_{p_i \in \mathcal{P}} P(p_i | T'_{p_i}) \quad (3)$$

Split	# claims	# languages
Train	19079	13
Development	2535	12
In-domain	3826	12
Out-of-domain	2368	4
Zero-shot	3381	12

Table 1: Dataset statistics of X-FACT.

Figure 2 demonstrates a graphical illustration.

However, if the translation is performed on every claim, the model would be discouraged from retrieving passages in the same language as the query, which is not a desirable property of the retriever. Therefore, we set an equal probability for a claim to be translated to any language or not being translated (i.e. the probability of not doing translation is $\frac{1}{7}$). We repurposed the passage collection as the training corpus for learning X-ICT.

3.2 Multilingual Reader

We use mBERT (Devlin et al., 2019), a multilingual version of BERT, as the encoder for our multilingual reader. The claim c and the corresponding metadata are encoded jointly with a template T : [Claim made by *Claimer* on *Claim-Date*, reported in *Language: Claim*], where *Claimer*, *Claim-Date*, and *Language* are placeholders. The claim template T and each of the k^2 retrieved passages p_i are first encoded independently

$$h_T = \text{mBERT}(T) [\text{CLS}] \quad (4)$$

$$h_{p_i} = \text{mBERT}(p_i) [\text{CLS}]. \quad (5)$$

The final prediction is then made by feeding the concatenated [CLS] embeddings into a multi-layer perceptron $\hat{y} = \text{MLP}([h_T; h_{p_0}; \dots; h_{p_k}])$. The model is optimized with the cross-entropy loss

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N y_i \log \hat{y}_i, \quad (6)$$

where y_i and \hat{y}_i denote the ground truth label and the predicted label of the i -th sample respectively, and N denotes the total number of samples.

4 Experimental Setup

4.1 Dataset and Evaluation Metric

Our experiments are conducted on a multilingual fact-checking dataset: X-FACT (Gupta and Srikumar, 2021). X-FACT contains 31,189 claims in 25 languages collected from fact-checking websites via Google’s Fact Check Explorer. Each

²Empirically, we set k to 5 for best overall performance.

claim is annotated with one of the following seven labels: TRUE, MOSTLY-TRUE, PARTLY-TRUE, MOSTLY-FALSE, FALSE, UNVERIFIABLE, and OTHER. Gupta and Srikumar (2021) split the data in a way that allows evaluation in various settings, as described in Table 1. *In-domain* and *out-of-domain* test sets contain claims in the same languages as those in the training set, except that the claims in the *out-of-domain* split are from different websites. Our main focus is the *zero-shot* setup, where there is no overlap between the languages in the *zero-shot* split and those in the training set. We use macro F1 as the evaluation metric, following Gupta and Srikumar (2021).

4.2 Baselines

We compare the following competitive retrieval systems using different retrieval components but with the same multilingual reader.

MT + DPR A common approach to cross-lingual tasks is translating inputs from target languages to source languages of richer resources so that stronger monolingual models can be utilized (Ahmad et al., 2021; Asai et al., 2021a). We translate all claims and all passages into English with the HELSINKI-NLP neural machine translation models³ for its comprehensive language coverage and decent performance. For languages not covered by HELSINKI-NLP, we use Google Translate instead. Then, we use DPR (Karpukhin et al., 2020) to perform retrieval based on the translated claims and passages.

BM25 BM25 (Robertson and Zaragoza, 2009) has demonstrated advantages over dense vector representation approaches in monolingual retrieval tasks (Lee et al., 2019). Since BM25 only works in a monolingual setup, we create dummy empty passages for claims whose languages are not presented in the passage collection P . Our implementation is based on the Rank-BM25 package⁴.

mDPR mDPR is a multilingual retriever based on DPR. It was trained on multilingual question answering datasets, as detailed in Section 3.

Google Search As demonstrated in previous work (Augenstein et al., 2019; Gupta and Srikumar, 2021), snippets from Google Search results can serve as evidence for fact-checking. We directly

³<https://huggingface.co/Helsinki-NLP>

⁴<https://pypi.org/project/rank-bm25/>

	Reader	Retrieval Method	Zero-shot F1 (%)	In-domain F1 (%)
Prior (Gupta and Srikumar, 2021)	Majority	None	7.6	6.9
	mBERT	None	16.7	39.4
	mBERT	Google Search	16.0	41.9
Ours	mBERT	None	17.25	36.91
	mBERT	Google Search	16.02	42.61
	mBERT	MT+DPR	15.01	35.29
	mBERT	BM25	17.43	38.29
	mBERT	mDPR	17.60	36.79
	mBERT	CONCRETE	19.83*	40.53

Table 2: Performance comparison in macro F1 (%) of various models on the X-FACT test sets. *None* retrieval method means not using any retrieval component, while *Majority* means predicting the majority label, FALSE, for all samples. When no retrieval method is used, the reader performs fact checking simply based on the claim template T described in Section 3. Statistical significance over the second best models computed using the paired bootstrap procedure (Berg-Kirkpatrick et al., 2012) is indicated with * ($p < .05$).

take the snippets obtained by Gupta and Srikumar (2021) as inputs.

4.3 Implementation Details

When trained with X-ICT on the passage collection, the retriever is optimized using AdamW (Loshchilov and Hutter, 2019) with a learning rate of $2e-5$ over 30 epochs. When fine-tuning the multilingual reader, we set the learning rate to $5e-5$ for parameters in mBERT and $1e-3$ for all other parameters. The maximum input sequence length for X-ICT and fine-tuning on X-FACT are set to 256 and 512, respectively. We use the pre-trained mBERT checkpoints on HuggingFace⁵.

5 Results

5.1 Main results

Table 2 summarizes the fact-checking performance on the X-FACT dataset. Our framework establishes a new state-of-the-art in zero-shot cross-lingual fact-checking, outperforming the previous best models by an absolute macro F1 of 2.23%. The improvements demonstrate the effectiveness of our approach in retrieving relevant passages from multiple languages to assist in fact-checking. Furthermore, we found that the use of vanilla mDPR does not improve performance compared to the absence of any retrieval component. This can be explained by the domain discrepancy issue for mDPR, as it was trained on question-like queries instead of claim-like queries. Furthermore, although CONCRETE is the most advantageous in the zero-shot setup, it trails behind Google Search in-domain setup. The significantly increased gap between in-domain F1 score and zero-shot F1 score

⁵<https://huggingface.co/bert-base-multilingual-cased>

for using Google Search suggests that the reader may exploit biases or patterns presented in Google Search’s results that are not transferrable across languages. To validate this hypothesis, we analyzed the relationship between the snippets returned by Google Search and the ground-truth labels. We found that for claims of richer resources in the training set, Google Search is able to retrieve evidence that strongly indicates the veracity of the claims due to the abundance of fact-checking websites. For example, among Indonesian claims where Google Search results contain the string “SALAH” (WRONG), 50% of them are PARTLY TRUE and 45% of them are FALSE. Such patterns can also be found in the *in-domain* split, but not in the *zero-shot* split. This finding explains the increased gap between the performance on these two splits when using Google Search as the retrieval method. It also implies that our approach is more generalizable to lower-resource languages and more applicable when no fact-checking websites have debunked the input claims.

5.2 Performance Analysis

Language Choice in Passage Collection

Throughout our experiments, we found that the language of the retrieved passage can affect the fact-checking performance. We hypothesize that it may be more challenging for readers to reason through passages when the passage language is distant from the claim language. To verify such a hypothesis, we first compute the distance between each pair of languages based on word ordering, following Ahmad et al. (2019). Then, we remove passages of a particular language from the passage collection. With this subset of selected passage collection, we use CONCRETE to retrieve relevant passages for fact-checking and




Figure 3: Performance difference in macro F1 when Indonesian passages are removed from the passage collection. On the x-axis, languages are sorted in ascending order by distance to Indonesian (i.e. Bengali is the closest, while Marathi is the furthest). We compute the distance between two languages based on word ordering, following [Ahmad et al. \(2019\)](#).

train another fact-checker based on newly retrieved passages. Finally, we compare the performance difference for different languages between this system and the model discussed in Section 3 in the *zero-shot* setup. Figure 3 shows the results for each language when Indonesian passages are removed from the passage collection. We observe that the performance drop and the distance between each language and Indonesian are highly correlated. In fact, we see that removing Indonesian is beneficial for some distant languages such as Dutch and Marathi.

Given these results, the following question arises: is this phenomenon caused by the poor capability of our multilingual reader to reason with passages and claims in distant languages, or by the higher information overlap between Indonesian passages and claims in closer languages? To better understand the results, we translate the retrieved passages into three languages of distinct language families: Indonesian, Portuguese, and Arabic. With the three sets of translated passages, we train three fact-checkers. Then, we compute the difference between the original performance and the performance achieved using the translated passages. In Figure 4, we observe that the performance difference is negatively correlated with the distance between the languages of the claim and the passage. The correlation coefficient is -0.490 per Kendall’s Tau ([Kendall, 1938](#)). This confirms that passages in distant languages are indeed harder for our multilingual reader to reason with. The trend is consistent with the findings of [Asai et al. \(2021b\)](#).

Impact of the Amount of Training Data To test the data efficiency of our approach, we compared CONCRETE with mDPR and Google Search




Figure 4: Performance difference in macro F1 when the retrieved passages are translated to Indonesian, Portuguese, and Arabic, respectively. Overall, the performance difference has a negative correlation with the distance between the passage’s language and the claim’s language.

in the *zero-shot* setup using different numbers of languages for training. As shown in Figure 5, CONCRETE consistently outperforms the other two methods across all settings. This indicates the strength of CONCRETE in aiding fact-checking in low- and high-resource scenarios.

Impact of Retrieving from Multiple Languages

We conducted a case study on cross-lingual retrieval versus monolingual retrieval to understand whether retrieving passages from multiple languages actually helps the performance. In particular, we train and evaluate the models on samples whose languages are in the passage collection⁶. For the monolingual retrieval setting, the model is restricted to retrieving passage in the same language as the input claim, while the cross-lingual setting does not have such a restriction. We found that for both mDPR and CONCRETE, cross-lingual retrieval setting outperforms their monolingual retrieval counterparts, as shown in Figure 6. This finding confirms that retrieving evidence in multiple languages helps cross-lingual transfer for fact-checking.

5.3 Qualitative Analysis

The following qualitative analysis provides an intuition for our model’s advantage in cross-lingual fact-checking.

Impact of X-ICT To validate the effectiveness of X-ICT in retrieving claims that are more relevant to the topic of the claim, we compared 50 predictions between models using CONCRETE and mDPR as retriever in the split *zero-shot*. The results show that 23 errors made by mDPR are corrected

⁶We train the models on Portuguese, Indonesian, and Arabic, and evaluate them on French, Persian, and Russian.

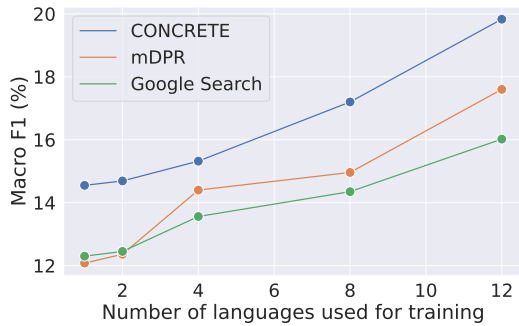


Figure 5: Zero-shot cross-lingual performance on X-FACT with regard to various numbers of languages used for training.

by CONCRETE, while only 6 new errors are introduced. We found that mDPR can often retrieve passages that are relevant to a part of a claim, but the topic of the retrieved passages may not align well with that of the claim, likely due to the domain mismatch issue discussed in Section 3. An example is shown in Figure 8. This reflects that X-ICT is able to improve the retrieval quality even though only pseudo feedback is used for training.

Importance of Passages’ Trustworthiness

Comparing the predictions and the retrieved passages between CONCRETE and Google Search on the *zero-shot* split, we observe that our approach is better at identifying FALSE and MOSTLY FALSE claims. As demonstrated in Figure 8, when a claim is FALSE or MOSTLY FALSE, the snippets returned by Google Search are often contradicted with each other, which usually leads to incorrect predictions. The inconsistency in the snippets from Google Search is caused by the fact that Google Search retrieves information from the entire Web without considering the trustworthiness of the source. On the contrary, our approach offers the flexibility to include only trustworthy information in the passage collection. In the *zero-shot* split, we found that this property of Google Search leads to 44 and 67 more errors in identifying FALSE and MOSTLY FALSE claims, respectively.

5.4 Remaining Challenges

To identify the remaining challenges, we compare 50 errors made by our model with ground-truth labels and analyze the sources of errors, as illustrated in Figure 7. The following paragraphs will discuss these categories with examples.

Evidence cannot be retrieved. The most common error is caused by the absence of supporting

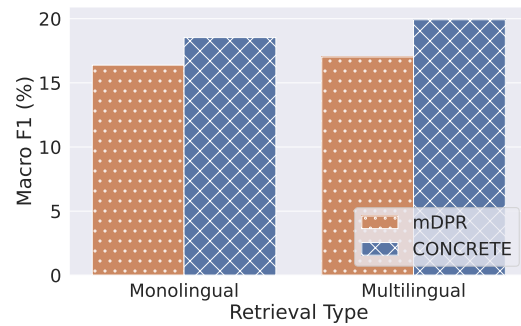


Figure 6: Performance comparison between using monolingual and multilingual retrieval on cross-lingual fact-checking. For both mDPR and CONCRETE, retrieving passages in multiple languages improves the overall performance.

or refuting evidence in the passage collection. For the majority of such errors, the claims are country specific. For example,

Mann med tre koner får tre leiligheter i Sverige. (Husband with three wives gets three apartments in Sweden.)

We can address this issue by adding trustworthy news articles from more countries into the passage collection.

Under-specified context. Another major source of errors is the underspecification of the input claim. An example claim is:

“Nuk besoj që janë të informuar as partnerët tanë, SHBA dhe NATO, sepse do të isha i informuar edhe unë.” (I do not believe that our partners, the US and NATO, are informed either, because I would be informed as well.)

In this claim, it does not specify who is “I” and what the US and NATO are not informed. Therefore, the given information is too little to determine the veracity of the claim. This problem could potentially be solved by mining the original context from the Internet.




Figure 7: Distribution of the remaining errors.

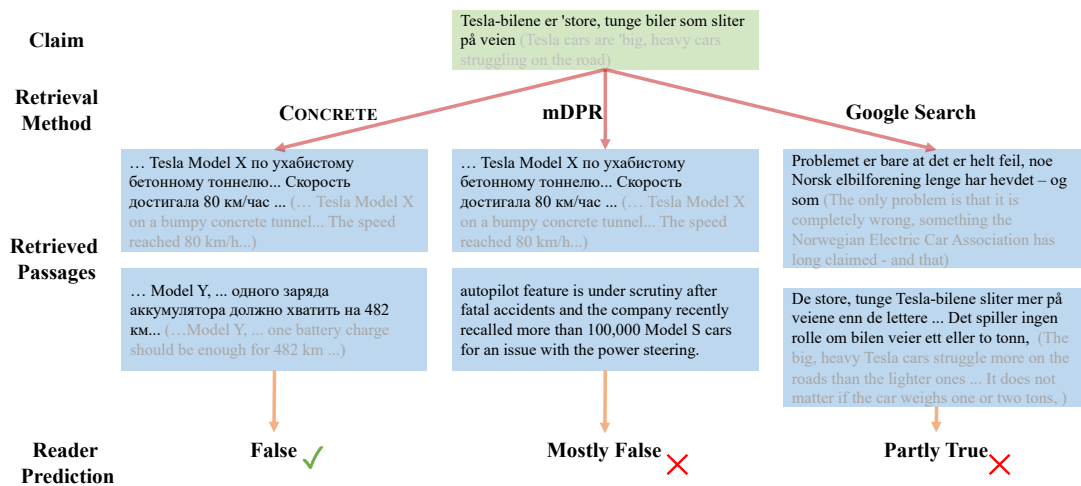


Figure 8: An example showing how relevant trustworthy passages retrieved by CONCRETE lead to correct predictions. CONCRETE retrieves two passages that prove Tesla cars are not slow or struggling on the road. Since the passages are from trustworthy sources, our reader can correctly predict the claim as FALSE. On the other hand, the second passage that mDPR retrieves is slightly related to the claim (Tesla) but not directly relevant to the topic (the poor performance of Tesla cars). For Google Search’s results, these two passages contradict each other as one is from a fact-checking website and the other is from an unreliable source. Hence, the model cannot predict correctly based on the retrieved passages from mDPR or Google Search.

Require intent identification. Some of the claims are correct but contain misleading information. For instance,

“Писатель Дин Кунц предсказал появление коронавируса в своей книге в 1981 году и называл его Ухань-400.” (Writer Dean Koontz predicted the emergence of the coronavirus in his book in 1981 and called it Wuhan-400.)

This claim is correct, but the claimer attempts to mislead the audience by linking the coronavirus in the book with COVID-19, which is false. To correctly predict this claim, the model should be capable of identifying the intent behind the claims.

Reader failure. In some cases, the reader fails to predict the correct veracity even though the supporting evidence is successfully retrieved. This is mainly due to the long distances between passages and claims, as discussed in Section 5.2.

Annotation error. The dataset was created by matching the rating of each claim on fact-checking websites with its label definition. Annotation errors could be caused by (1) the misalignment between the rating and the label definition, and (2) the inaccurate ratings listed on the fact-checking websites. We found that each case accounts for about half of the annotation errors.

6 Related Work

6.1 Fact-checking

Previous fact-checking approaches can be roughly divided into two categories based on task formulations. The first type of formulation assumes that evidence candidates are given, such as the FEVER dataset (Thorne et al., 2018) and the SCIFACT dataset (Wadden et al., 2020). Previous approaches for this category of fact-checking tasks often involve a retrieval module to retrieve relevant evidence from the given candidate pool followed by a reasoning component that determines the compatibility between a piece of evidence and the input claim (Yin and Roth, 2018; Pradeep et al., 2021). The second category is the *open-retrieval* setting⁷, where evidence candidates are not provided, such as the LIAR dataset (Wang, 2017) and the X-FACT dataset (Gupta and Sriku-mar, 2021). For this task formulation, one of the main challenges is where and how to retrieve evidence. Some work determines the veracity of a claim based solely on the claim itself and the information learned by language models during the pre-training stage (Lee et al., 2021). However, such an approach is tied to the period of time in which the pre-training data is collected and does not generalize well to new claims. Other studies

⁷We borrow the term *open-retrieval* from the field of question answering.

devise hand-crafted linguistic features as input to the fact-checking models (Mihalcea and Strapparava, 2009; Choudhary and Arora, 2021). However, these approaches are language-specific, whereas our approach is language-agnostic since it consists of a cross-lingual retriever and a multilingual language model. Gupta and Srikumar (2021) use a similar approach, which retrieves relevant snippets using Google Search instead of a cross-lingual retriever. Our experimental results show that the proposed cross-lingual retriever is more effective than Google Search in the *zero-shot* setting due to the fact that Google Search ignores the trustworthiness of the retrieved information and that downstream models tend to exploit the biased patterns in Google Search results that are not transferrable across languages, as shown in Section 5.

6.2 Cross-lingual Retrieval

Early attempts on cross-lingual retrieval adopt a pipeline consisting of a statistical machine translation system and a monolingual retrieval model (Hiemstra and de Jong, 1999; Ture and Lin, 2013). These methods do not perform well due to the poor performance of statistical machine translation systems. Later work addresses this issue with bilingual embeddings (Vulic and Moens, 2015; Litschko et al., 2018). More recently, large pre-trained multilingual language models demonstrate significant advantages in constructing multilingual representations (Jiang et al., 2020). Yu et al. (2021) pre-train a cross-lingual language model tailored for the retrieval tasks. Yet, the computation complexity is relatively high due to the cross-encoder architecture of the model. Namely, it takes a pair of query and evidence as inputs, instead of encoding the query and evidence independently. The mDPR model presented in (Asai et al., 2021b) is the most favorable for our task due to its high efficiency and performance. However, mDPR was trained on datasets where queries are questions instead of claims. Therefore, domain adaptation is needed for mDPR to be applied to our task. CONCRETE extends mDPR by adapting it to fact-checking using a self-supervised cross-lingual retrieval algorithm to mitigate the domain discrepancy problem while maintaining high efficiency.

7 Conclusions and Future Work

We have proposed CONCRETE, a claim-oriented cross-lingual retriever that retrieves trustworthy

passages from a multilingual passage collection. To overcome the lack of IR training data with claim-like queries, we present the Cross-lingual Inverse Cloze Task (X-ICT) that leverages pseudo feedback to train the retriever. Experimental results on X-FACT showed that our approach outperforms all previous systems in the *zero-shot* cross-lingual setting. For future work, we plan to investigate the adaptive selection mechanism for passages based on distances and develop more robust readers for reasoning through passages of longer distances via representation learning.

8 Ethical Considerations

Although our framework has significant advantages over the previous state of the art, the proposed model is still far from being a reliable cross-lingual fact checker given its great potential for improvement. If such a system is deployed for public use, the general public could lose trust in automatic fact-checking systems, and the situation of infodemic can exacerbate. Therefore, at the current stage, our system should be served as an assistant for human fact-checkers to validate the veracity of claims instead of directly applying for public use, especially in the *zero-shot* setting. With our framework, the efficiency of manual fact-checking can be significantly improved thanks to its ability to retrieve relevant information across multiple languages and produce a reasonably good preliminary judgement on the veracity of the input claim.

In addition, we also acknowledge that the use of large multilingual language models pre-trained on the Web could lead to biased outputs. Fortunately, after a close inspection into the X-FACT dataset, we do not find such biased patterns in it. This means that fine-tuning our proposed framework on X-FACT should alleviate the problem of biased predictions.

Acknowledgement

This research is based upon work supported by U.S. DARPA SemaFor Program No. HR001120C0123. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Wasi Ahmad, Zhisong Zhang, Xuezhe Ma, Eduard Hovy, Kai-Wei Chang, and Nanyun Peng. 2019. [On difficulties of cross-lingual transfer with order differences: A case study on dependency parsing](#). 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 2440–2452, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wasi Uddin Ahmad, Nanyun Peng, and Kai-Wei Chang. 2021. [Gate: Graph attention transformer encoder for cross-lingual relation and event extraction](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Akari Asai, Jungo Kasai, Jonathan Clark, Kenton Lee, Eunsol Choi, and Hannaneh Hajishirzi. 2021a. [XOR QA: Cross-lingual open-retrieval question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 547–564, Online. Association for Computational Linguistics.
- Akari Asai, Xinyan Yu, Jungo Kasai, and Hannaneh Hajishirzi. 2021b. [One question answering model for many languages with cross-lingual dense passage retrieval](#). In *NeurIPS*.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. [MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims](#). 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 4685–4697, Hong Kong, China. Association for Computational Linguistics.
- Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. [An empirical investigation of statistical significance in NLP](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 995–1005, Jeju Island, Korea. Association for Computational Linguistics.
- Anshika Choudhary and Anuja Arora. 2021. [Linguistic feature based learning model for fake news detection and classification](#). *Expert Systems with Applications*, 169:114171.
- 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.
- Ogi Djuraskovic. 2022. [47 blogging statistics \(2022\): How many blogs are there?](#)
- Yi Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal, and Avi Sil. 2021. [InfoSurgeon: Cross-media fine-grained information consistency checking for fake news detection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1683–1698, Online. Association for Computational Linguistics.
- Yi R. Fung, Kung-Hsiang Huang, Preslav Nakov, and Heng Ji. 2022. [The battlefield of combating misinformation and coping with media bias](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*, page 4790–4791, New York, NY, USA. Association for Computing Machinery.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. [A Survey on Automated Fact-Checking](#). *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Ashim Gupta and Vivek Srikumar. 2021. [X-factor: A new benchmark dataset for multilingual fact checking](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 675–682, Online. Association for Computational Linguistics.
- Djoerd Hiemstra and Franciska de Jong. 1999. [Disambiguation Strategies for Cross-Language Information Retrieval](#). In *Proceedings of the Third European Conference on Research and Advanced Technology for Digital Libraries*, volume 1696 of *Lecture Notes in Computer Science*, pages 274–293. Springer.
- Thi Bich Ngoc Hoang and Josiane Mothe. 2018. [Location extraction from tweets](#). *Information Processing & Management*, 54(2):129–144.
- Kung-Hsiang Huang, Kathleen McKeown, Preslav Nakov, Yejin Choi, and Heng Ji. 2022. [Faking fake news for real fake news detection: Propaganda-loaded training data generation](#). *arXiv preprint arXiv:2203.05386*.
- Zhuolin Jiang, Amro El-Jaroudi, William Hartmann, Damianos Karakos, and Lingjun Zhao. 2020. [Cross-lingual information retrieval with BERT](#). In *Proceedings of the workshop on Cross-Language Search and Summarization of Text and Speech (CLSSTS2020)*, pages 26–31, Marseille, France. European Language Resources Association.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural*

- Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Ashkan Kazemi, Kiran Garimella, Devin Gaffney, and Scott Hale. 2021. [Claim matching beyond English to scale global fact-checking](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4504–4517, Online. Association for Computational Linguistics.
- Ashkan Kazemi, Zehua Li, Verónica Pérez-Rosas, Scott A Hale, and Rada Mihalcea. 2022. Matching tweets with applicable fact-checks across languages. *arXiv preprint arXiv:2202.07094*.
- Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. 2021. [Towards few-shot fact-checking via perplexity](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1971–1981, Online. Association for Computational Linguistics.
- Robert Litschko, Goran Glavas, Simone Paolo Ponzetto, and Ivan Vulic. 2018. Unsupervised cross-lingual information retrieval using monolingual data only. *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Rada Mihalcea and Carlo Strapparava. 2009. [The lie detector: Explorations in the automatic recognition of deceptive language](#). In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pages 309–312, Suntec, Singapore. Association for Computational Linguistics.
- Umaru A Pate, Danjuma Gambo, Adamkolo Mohammed Ibrahim, et al. 2019. The impact of fake news and the emerging post-truth political era on nigerian polity: A review of literature. *Studies in Media and Communication*, 7(1):21–29.
- Ronak Pradeep, Xueguang Ma, Rodrigo Nogueira, and Jimmy Lin. 2021. [Scientific claim verification with VerT5erini](#). In *Proceedings of the 12th International Workshop on Health Text Mining and Information Analysis*, pages 94–103, online. Association for Computational Linguistics.
- Stephen Robertson and Hugo Zaragoza. 2009. *The probabilistic relevance framework: BM25 and beyond*. Now Publishers Inc.
- Dietram A Scheufele and Nicole M Krause. 2019. Science audiences, misinformation, and fake news. *Proceedings of the National Academy of Sciences*, 116(16):7662–7669.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#).
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). 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 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Ferhan Ture and Jimmy Lin. 2013. Flat vs. hierarchical phrase-based translation models for cross-language information retrieval. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 813–816.
- Ivan Vulic and Marie-Francine Moens. 2015. Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or fiction: Verifying scientific claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics.
- 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 (Volume 2: Short Papers)*, pages 422–426, Vancouver, Canada. Association for Computational Linguistics.
- Xueqing Wu, Kung-Hsiang Huang, Yi Fung, and Heng Ji. 2022. [Cross-document misinformation detection based on event graph reasoning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 543–558, Seattle, United States. Association for Computational Linguistics.
- Wenpeng Yin and Dan Roth. 2018. [TwoWingOS: A two-wing optimization strategy for evidential claim verification](#). In *Proceedings of the 2018 Conference*

on Empirical Methods in Natural Language Processing, pages 105–114, Brussels, Belgium. Association for Computational Linguistics.

Puxuan Yu, Hongliang Fei, and P. Li. 2021. Cross-lingual language model pretraining for retrieval. *Proceedings of the Web Conference 2021*.