

ViWikiFC: Fact-Checking for Vietnamese Wikipedia-Based Textual Knowledge Source

Hung Tuan Le, Long Truong To, Manh Trong Nguyen, Kiet Van Nguyen*

University of Information Technology, Ho Chi Minh City, Vietnam

Vietnam National University, Ho Chi Minh City, Vietnam

{21520250, 21521101, 21520343}@gm.uit.edu.vn, kietnv@uit.edu.vn

Abstract

Fact-checking is essential due to the explosion of misinformation in the media ecosystem. Although false information exists in every language and country, most research to solve the problem has mainly concentrated on huge communities like English and Chinese. Low-resource languages like Vietnamese are necessary to explore corpora and models for fact verification. To bridge this gap, we construct ViWikiFC, the first manually annotated open-domain corpus for Vietnamese Wikipedia Fact Checking more than 20K claims generated by converting evidence sentences extracted from Wikipedia articles. We analyze our corpus through many linguistic aspects, from the new dependency rate, the new n-gram rate, and the new word rate. We conducted various experiments for Vietnamese fact-checking, including evidence retrieval and verdict prediction. BM25 and InfoXLM_{Large} achieved the best results in two tasks, with BM25 achieving an accuracy of 88.30% for SUPPORTS, 86.93% for REFUTES, and only 56.67% for the NEI label in the evidence retrieval task. InfoXLM_{Large} achieved an F₁ score of 86.51%. Furthermore, we also conducted a pipeline approach, which only achieved a strict accuracy of 67.00% when using InfoXLM_{Large} and BM25. These results demonstrate that our dataset is challenging for the Vietnamese language model in fact-checking tasks.

Keywords: Fact Checking, Language Model, Information Verification, Corpus

1. Introduction

Following the increase in the amount of information, the lack of strict policies when spreading information has led to misinformation and disinformation on social media. This can cause conflict and manipulate the group of people. To minimize these impacts, organizations like PolitiFact and FactCheck.org play an essential role in online Fact Verification by manually verifying claims based on different sources of evidence. However, manual verification is time-consuming and insufficient given the speed of information updates on social media.

Automatic Fact-Checking is a complicated task that can be separated into four sub-tasks: claim detection, evidence retrieval, verdict prediction, and justification production (Guo et al., 2022). While there are many works on improving fact-checking systems (Lewis et al., 2020; Maillard et al., 2023; Pan et al., 2023) and large-scale corpora for various languages such as NELA (Horne et al., 2018), FakeCovid (Shahi and Nandini, 2020), FEVER (Thorne et al., 2018), and VitaminC (Schuster et al., 2021), Vietnamese remains a low-resource language for NLP, especially in automatic fact-checking.

To contribute to Vietnamese NLP research, we present ViWikiFC: the first large-scale, open-domain corpus for Vietnamese Fact-Checking on Wikipedia. The corpus consists of 20,916 manually annotated claims based on evidence from

Wikipedia pages. Following FEVER (Thorne et al., 2018), our corpus has three label classes: SUPPORTS, REFUTES, and NOTENOUGHINFORMATION (NEI). The evidence is manually rewritten into three types of claims, making them more realistic and semantically diverse than FEVER’s approach. We also follow ViNLI (Huynh et al., 2022) to create double claims with different meanings for each label from one evidence.

Annotators who are native Vietnamese speakers with education beyond high school were carefully trained to ensure consistency. We conducted experiments on evidence retrieval using TF-IDF, BM25, and Vietnamese-SBERT (Phan et al., 2022b), and on verdict prediction using CBOW (Mikolov et al., 2013), BiLSTM (Graves and Schmidhuber, 2005), multilingual models (mBERT (Kenton and Toutanova, 2019), XLM-R (Conneau et al., 2020), InfoXLM (Chi et al., 2021)), and Vietnamese-specific models (PhoBERT (Nguyen and Tuan Nguyen, 2020), ViDeBERTa (Tran et al., 2023)).

Contributions in this paper are described as follows.

- Firstly, we propose the first fact-checking corpus on Vietnamese Wikipedia, comprising 20,916 claim sentences manually annotated based on 73 Wikipedia articles built with a consensus agreement among annotators, achieving a 95.87% Fleiss’k-agreement.
- Second, we conduct two experiments on two tasks, evidence retrieval and verdict predic-

*Corresponding author.

Corpus	Claims	Language	Source	Domain	Annotated Evidence	Annotated Label
PolitiFact (2014)	106	English	PolitiFact	Politics	✗	✗
Liar (2017)	12,836	English	PolitiFact	Politics	✗	✗
FEVER (2018)	185,445	English	Wiki	Multiple	✓	✓
PUBHEALTH (2020)	11,832	English	Fact-Checking/News	Health	✗	✗
HOVER (2020)	26,171	English	Wiki	Multiple	✓	✓
TabFact (2020)	92,283	English	Wiki	Multiple	✗	✓
InfoTabs (2020)	23,738	English	Wiki	Multiple	✗	✓
ANT (2020)	4,547	Arabic	News	Multiple	✗	✗
FakeCovid (2020)	5,182	Multilingual (3)	Fact Check	Health	✗	✗
X-Fact (2021)	31,189	Multilingual (25)	Fact Check	Multiple	✓	✗
VitaminC (2021)	488,904	English	Wiki	Multiple	✗	✓
DanFEVER (2021)	6,407	Danish	Wiki	Multiple	✓	✓
CHEF (2022)	10,000	Chinese	Internet	Multiple	✓	✓
ViWikiFC	20,916	Vietnamese	Wiki	Multiple	✗	✓

Table 1: Overall statistics of fact-checking corpora.

tion, while also developing a pipeline approach for fact-checking models, encompassing both neural network-based and pre-trained transformer-based models.

- Next, we analyze the corpus from various linguistic perspectives in the Vietnamese language to gain more insight into Vietnamese fact-checking, including the new n-gram rate, the new dependency rate, and the new word rate.
- Finally, our corpus and annotation tool are available for research purposes.

2. Related Work

In this section, we review prior studies and resources related to fact-checking, focusing on corpora and baseline models that have shaped this research direction. Section 2.1 introduces key datasets, while Section 2.2 presents the main methodological approaches.

2.1. Fact-Checking Corpora

We reviewed existing fact-checking datasets, summarized in Table 1, categorized by dataset size, language, source, domain, and annotation types.

The PolitiFact corpus (Vlachos and Riedel, 2014) laid the foundation for modern fact-checking by defining three core tasks: claim detection, evidence retrieval, and verdict prediction. However, its limited size restricted model training. The LIAR dataset (Wang, 2017) expanded on this with over 12.8K samples from PolitiFact¹ but lacked evidence annotations. FEVER (Thorne et al., 2018) later introduced a large-scale corpus of 180K manually written claims based on Wikipedia², establishing an automated pipeline structure for verification. Subsequent corpora such as Snopes (Hanselowski

et al., 2019) and MultiFC (Augenstein et al., 2019) diversified sources and label types, while later datasets like HoVer (Jiang et al., 2020) and VitaminC (Schuster et al., 2021) expanded evidence coverage and reasoning complexity.

For semi-structured data, TabFact (Chen et al., 2020), InfoTabs (Gupta et al., 2020), and FEVEROUS (Aly et al., 2021) introduced fact verification using tables or hybrid text-table formats. Multilingual efforts have also emerged, including X-Fact (Gupta and Srikumar, 2021), ANT (Khouja, 2020), DanFEVER (Nørregaard and Derczynski, 2021), and CHEF (Hu et al., 2022).

Despite this progress, Vietnamese NLP still lacks an open-domain, large-scale, and well-annotated fact-checking corpus. To address this gap, we constructed a Vietnamese fact-checking dataset and developed an accompanying verification pipeline using pre-trained Vietnamese language models.

2.2. Fact-Checking Methods

Fact-checking typically involves four sub-tasks: claim detection, evidence retrieval, verdict prediction, and justification generation. Among these, evidence retrieval and verdict prediction are central, forming the core of our research.

Evidence retrieval involves finding relevant information from various sources to support or refute a claim. Traditional non-neural methods such as TF-IDF (Thorne et al., 2018; Hu et al., 2022) and BM25 (Chen et al., 2017) rely on lexical matching but are limited in semantic understanding. Neural approaches, such as Dense Passage Retrieval (DPR) (Karpukhin et al., 2020), use dense vector representations via dual-encoder architectures, achieving significant improvements in open-domain QA (Karpukhin et al., 2020; Chang et al., 2019; Lee et al., 2019) and fact-checking (Samarinas et al., 2021, 2020). Cross-encoder re-rankers based on BERT and its variants (Nogueira and Cho, 2019; Nogueira et al., 2019; Wang et al., 2019) further enhance retrieval precision.

¹<https://www.politifact.com/>

²<https://www.wikipedia.org/>

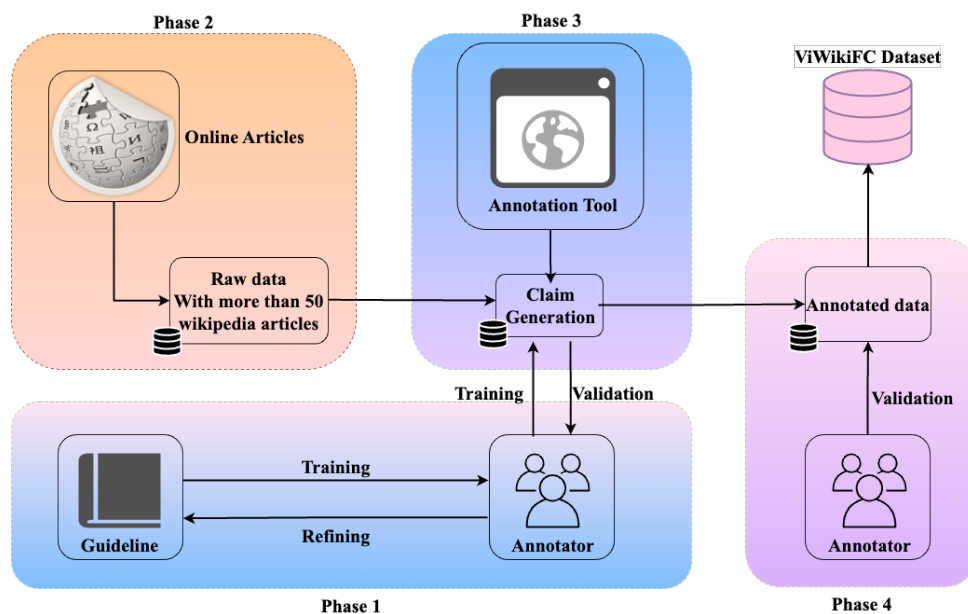


Figure 1: The creation process of the fact-checking corpus.

Verdict prediction determines the truthfulness of a claim based on retrieved evidence. It is typically formulated as a sentence-pair classification problem. Early approaches used CBOW (Mikolov et al., 2013), RNN (Elman, 1990), or BiLSTM (Graves and Schmidhuber, 2005), while transformer-based models such as BERT (Devlin et al., 2019) and its successors—RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2019), and XLM-R (Conneau et al., 2020)—greatly improved performance. In Vietnamese, PhoBERT (Nguyen and Tuan Nguyen, 2020) was the first monolingual pre-trained model, trained on a 20GB corpus of Wikipedia and news data, followed by ViT5 (Phan et al., 2022a) and ViDeBERTa (Tran et al., 2023), which further enhanced representation quality for downstream tasks, including fact-checking.

3. Corpus Creation

The corpus was created throughout four phases (see Figure 1), including annotator recruitment and training (see Section 3.1), evidence selection (see Section 3.2), claim generation (see Section 3.3), and corpus validation (see Section 3.4). Furthermore, we also analyzed the corpus (see Section 3.5) in terms of various linguistic aspects to gain a deeper understanding of the language characteristics of the corpus.

3.1. Annotator Recruitment and Training

We recruited thirty Vietnamese native speakers with at least high school to college literacy to create claims based on given evidence sentences. All annotators were familiar with online reading and

information searching. Before participating in the official annotation, they underwent a two-phase training process to ensure annotation quality, with agreement measured by the Fleiss κ score (Fleiss, 1971).

In the first phase, annotators were trained using detailed guidelines to understand claim-generation rules. They then annotated a small training corpus containing 50 evidence sentences labeled as *SUPPORTS*, *REFUTES*, or *NOT ENOUGH INFORMATION*. Each claim-evidence pair was reviewed by one of three authors, and inter-author agreement above 0.90 was required for the annotator to proceed to the second phase. Annotators scoring below 0.90 were asked to review their mistakes and repeat the first phase using a new training set.

In the second phase, the thirty annotators were randomly divided into six groups of five. Each annotator labeled at least one dataset folder of 150 claim-evidence pairs, with labels concealed. Every group then re-labeled five folders produced by other groups, and inter-annotator agreement was again evaluated using Fleiss κ . Annotators achieving a consensus score of 0.95 or higher qualified for the official annotation process; otherwise, they repeated the first training phase. This procedure continued until all annotators achieved the required agreement threshold. During pilot annotations, annotators were encouraged to document ambiguous cases (e.g., polysemy, unclear entities) and note common labeling errors to refine the annotation guidelines.

3.2. Evidence Selection

We use Wikipedia as the primary data source to build the ViWikiFC corpus. We do not prioritize establishing the credibility of the data source, as, in reality, all data sources can sometimes present inaccuracies or omit information, either unintentionally or deliberately. Therefore, no source can be deemed completely trustworthy. The team considers the coverage of information in various aspects of life as a crucial factor. Wikipedia, being an open encyclopedia with more than 55 million articles on a wide array of topics, helps to build a corpus with greater comprehensiveness compared to constructing it solely from other sources, such as political information from Politifact³ or news articles from VnExpress⁴. We extracted 3,812 evidence sentences from more than 1,479 paragraphs in 73 articles on Wikipedia, covering various topics such as history, geography, philosophy, and science.

3.3. Claim Generation

Following the FEVER dataset (Thorne et al., 2018), we design three labels—SUPPORTS, REFUTES, and NOT ENOUGH INFORMATION (NEI)—to capture the different relationships between a claim and its supporting evidence. In practice, many claims cannot be conclusively verified or falsified due to incomplete or ambiguous information; therefore, we include the NEI label to better reflect these real-world scenarios. Annotators are instructed to compose claim sentences for each of the three labels as follows:

- **SUPPORTS:** The claim can be determined as true based solely on the given evidence.
- **REFUTES:** The claim can be determined as false based solely on the given evidence.
- **NOT ENOUGH INFORMATION (NEI):** The claim cannot be judged as true or false using the provided evidence.

For the SUPPORTS and REFUTES labels, annotators write claims that rely only on the content of the evidence sentence, without using external knowledge or personal understanding. They are encouraged to reformulate, restructure, or vary the expression to enhance linguistic diversity and avoid simply paraphrasing the evidence. This approach differs from FEVER (Thorne et al., 2018), where claims often derive from minimal factual transformations.

For the NEI label, annotators must still reuse at least one component from the evidence (subject,

object, or event) while introducing new but contextually relevant information not present in the original sentence. This design ensures the NEI label remains semantically related to the evidence and also enables a more realistic evaluation of evidence retrieval and claim verification models.

Each evidence sentence leads to the creation of two distinct claim sentences for each label, resulting in a total of six claims per evidence. This procedure follows the approach in OCNLI (Hu et al., 2020) and IndoNLI (Mahendra et al., 2021), where annotators are instructed to explore multiple information pieces within a single evidence and use diverse linguistic forms to represent similar meanings. Furthermore, for a given evidence sentence, the claims corresponding to different labels show clear variation in writing style, word choice, and focus, thereby improving the dataset’s semantic diversity.

During annotator training, we observed that a noticeable portion of annotators struggled with constructing SUPPORTS claims, especially when the evidence sentence contained limited information (typically 15–25 tokens). It was challenging for them to compose a natural claim without overly mirroring the evidence. Similarly, the REFUTES label posed difficulties due to excessive or unclear negation, which occasionally caused label drift toward NEI.

3.4. Corpus Validation

To ensure data quality and annotation consistency, we perform a two-stage validation process: annotator validation and author validation, both conducted during the claim generation phase.

3.4.1. Annotator Validation

When 50% of the corpus was completed, the top five annotators were selected to relabel a separate subset of 750 claim–evidence pairs that none had previously worked on. Labels were re-assigned collaboratively, and pairs with fewer than three matching labels were discarded. The inter-annotator agreement reached 99.87% when at least three annotators agreed and 95.47% when at least two agreed. The Fleiss κ score was 95.87%, significantly higher than FEVER’s 68.41%, confirming strong consistency among annotators.

3.4.2. Author Validation

The same 750 pairs were further reviewed by the authors to assess guideline compliance. Results show that 96.12% of the claims strictly followed the annotation rules, while only 3.87% contained minor writing or labeling errors, demonstrating the overall reliability of the corpus.

³<https://www.politifact.com/>

⁴<https://vnexpress.net/>

3.5. Corpus Analysis

3.5.1. Overall Statistics

We divide the corpus, which consists of 20,916 claim sentences, randomly into three different sets: 80.00% for training (Train), 10.00% for development (Dev), and 10.00% for the test (Test) set for the VP task. For ER tasks, we do not employ pre-trained models on a portion of the data and then evaluate the remaining part. Instead, we use the cosine similarity measure between sentences within the model to extract evidence sentences, and this process will be carried out and evaluated on the entire corpus. The distribution of labels is summarized in Table 2.

	SUPPORTS	REFUTES	NEI
Train	5,594	5,573	5,571
Dev	666	694	730
Test	708	706	677

Table 2: Overall statistics of our corpus.

3.5.2. Length Distribution

Allocate the sentences of claims and evidence based on their length, as presented in Figure 3 and Figure 2. The length of claims is shorter compared to the evidence, but they still ensure a complete and accurate representation of the information conveyed by the evidence sentences. In reality, claim sentences often tend to be shorter in length, whereas evidence sentences are typically longer, serving the role of providing additional details and supporting the claim. Furthermore, writing too long sentences can result in providing more context and information to the model, making training and prediction less challenging and more straightforward. The shortest length for a claim is four words, while the longest when generating is 113 words. Sentences within the length range of 12 to 25 words constitute the largest proportion of the corpus.

3.5.3. New Word Rate

To evaluate diversity in the corpus, we measure the rate of new words in the claim sentence that do not appear in the evidence sentence. We use VnCoreNLP (Vu et al., 2018) for word segmentation. The results in table 3 show that the new word rate in the REFUTES label is the lowest, reaching only 31.28%. The word diversity of the NEI label (50.44%) is the highest among the three labels, followed by the SUPPORTS label with 32.91%. Both the SUPPORTS and REFUTES labels have a low new word rate compared to the NEI label, indicating that annotators tend to use fewer new words when constructing sentences for these two labels.

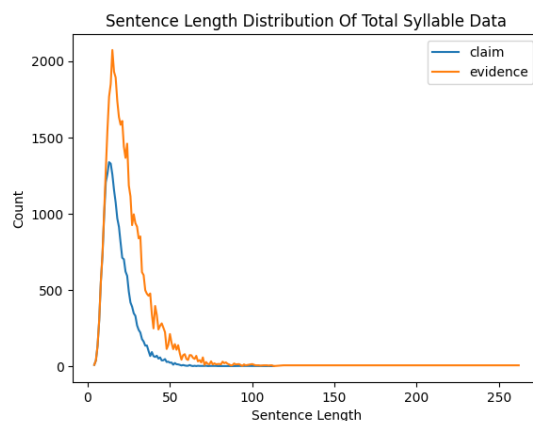


Figure 2: The distribution of total syllables in the corpus.

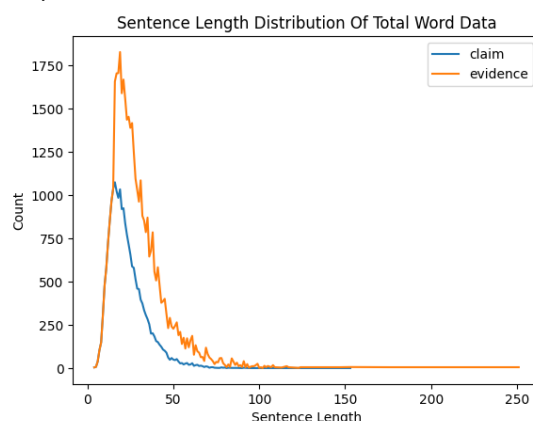


Figure 3: The distribution of total words in the corpus.

This tendency ensures semantic accuracy with the evidence sentence when external knowledge is not used. In Addition, we utilized part-of-speech analysis on new words to delve deeper into annotator data construction trends. We used PhoNLP (Nguyen and Nguyen, 2021) for word segmentation and summarized the results in Table 3. The statistical results reveal that nouns and verbs are the two primary components annotators use when composing claim sentences.

3.5.4. Corpus-generation rules analysis

To evaluate and analyze the linguistic factors of the annotators during the creation of the corpus, we proceed with the analysis of the rules used based on the ViNLI (Huynh et al., 2022) corpus generation process. We randomly chose 500 SUPPORTS and 500 REFUTES claim-evidence pairs from the validation set. Selecting 500 sentences for each label will increase the coverage of the data compared to analyzing a smaller sample of only 100 sentences, as was performed in ViNLI. The statistical results and illustrative examples are provided in Appendix

Label	New Word Rate (%)	New Dependency Rate (%)	Part-Of-Speech (%)					
			Noun	Verb	Adjective	Preposition	Adjunct	Other
SUPPORTS	32.91	75.26	26.97	30.22	7.65	10.81	7.92	16.45
REFUTES	31.28	67.90	26.27	25.47	8.66	8.04	13.10	18.47
NEI	50.44	81.96	33.96	23.61	8.96	9.82	7.65	15.99

Table 3: Corpus analysis in terms of linguistic aspects.

A.

For the SUPPORTS label, the rule annotator uses most is 'replace words with synonyms' with 58.60%; meanwhile, the 'turn adjective into relative clause' rule has the lowest ratio (Table 7). As for the REFUTES label, annotators tend to generate claims with the 'wrong reasoning about an event' rule, whereas 'opposite of time' and 'create a sentence that has the opposite meaning of a pre-supposition' share the same proportion of 5.60%, making them the least common rule used when creating data (Table 8).




Figure 4: The ratio of SUPPORT claim generation rules.




Figure 5: The ratio of REFUTE claim generation rules.

During the corpus generation process, annotators can use more than one rule for every claim they write. The analysis we conducted on claims shows that almost half of SUPPORT claims are written using two rules (49.80%), while one rule has 41.20%. For the REFUTES label, claims are nearly generated by one rule with a proportion of 88.00% (Figure 5). This can be explained by the fact that when using too many rules without reviewing carefully, REFUTES claims can be turned into NEI labels.

4. Empirical Evaluation

4.1. Baseline Models and Pipeline Description

To comprehensively evaluate the ViWikiFC dataset, we define two main subtasks: Evidence Retrieval (ER) and Verdict Prediction (VP), which are further integrated into a complete fact-checking pipeline.

4.1.1. Evidence Retrieval

In the ER task, given a claim sentence, the objective is to identify the most relevant evidence sentences from the corpus. We employ three retrieval methods: TF-IDF, Okapi BM25, and Vietnamese Sentence-BERT (SBERT). The TF-IDF model is implemented from scratch, BM25 is applied using the Rank-BM25⁵ library with default hyperparameters, and Vietnamese Sentence-BERT is adopted to capture semantic similarity between claims and evidence through sentence embeddings. This model, based on the Sentence-BERT architecture, enables efficient comparison of sentence-level meanings with reduced computational cost compared to traditional transformer-based encoders.

4.1.2. Verdict Prediction

For the VP task, we explore a range of models from traditional neural architectures to pre-trained language models. Simpler models, such as Continuous Bag of Words (CBOW) and BiLSTM, are trained with pre-trained PhoW2V embeddings at both syllable and word levels. We further investigate the performance of advanced transformer-based models, including Multilingual BERT (mBERT), XLM-RoBERTa, InfoXLM, PhoBERT, and ViDeBERTa, covering both multilingual and monolingual representations for Vietnamese. Word-based models (PhoBERT, ViDeBERTa) employ the VNCORENLP toolkit for word segmentation, while syllable-based models (mBERT, XLM-R, InfoXLM) use raw tokenized text. To ensure robustness, each experiment is conducted five times with different random seeds, and the standard deviation of the performance metrics is reported.

⁵<https://pypi.org/project/rank-bm25/>

4.1.3. Pipeline Description

Finally, we selected the models that achieved the highest performance in both tasks, evidence retrieval and verdict prediction, and integrated them into a complete fact-checking pipeline. This pipeline takes a claim sentence as input and outputs both an evidence sentence extracted from the corpus and a corresponding verdict label that indicates whether the claim is *SUPPORTED*, *REFUTED*, or *NOT ENOUGH INFORMATION*.

To ensure rigorous and reliable evaluation, we apply a **strict accuracy** metric, which considers a prediction correct only when both the retrieved evidence and the predicted label are accurate. In other words, if either the evidence or the verdict is incorrect, the entire prediction is counted as incorrect. This metric reflects the end-to-end quality of the fact-checking process rather than evaluating each component in isolation.

Strict accuracy is formally defined as follows, where v and v' denote the ground-truth and predicted evidence in the evidence retrieval task, and e and e' represent the ground-truth and predicted labels in the verdict prediction task $e, e' \in \{SUPPORTS, REFUTES, NOTENOUGHINFORMATION\}$:

$$\text{Strict Accuracy} = \delta(v, v') \times \delta(e, e')$$

where $\delta(a, b)$ equals 1 if $a = b$ and 0 otherwise.

4.1.4. Evaluation Metrics

To assess the performance of the models in these tasks, we employ accuracy as the primary evaluation metric. Additionally, we utilize the F_1 score (macro average) as the secondary evaluation metric for the verdict prediction task (VP), which is described below:

$$F_1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

4.2. Experimental Results

4.2.1. Evidence Retrieval

The experimental results presented in Table 5 shed light on the performance of models in the evidence retrieval (ER) task with pre-trained model Vietnamese-SBERT (SBERT), TF-IDF, and BM25. These results offer valuable information on the strengths and weaknesses of these approaches.

We can observe that BM25 consistently outperforms TF-IDF and SBERT across all 3 labels in both single-sentence and 5-sentence retrieval scenarios. As the number of retrieved sentences increases from 1 to 5, the accuracy of all 3 models increases, with BM25 achieving the highest accuracy, reaching 93.93% for the SUPPORTS label and 93.04% for the REFUTES label. SBERT, with

its transformer architecture, possesses the ability to capture the contextual nuances of the sentences it retrieves, but in the ViWikiFC corpus, the model's performance underperforms BM25. This suggests that BM25, despite its less complex architecture, exhibits remarkable performance in evidence retrieval tasks compared to other transformer-based methods (Thakur et al., 2021).

However, it is important to address the relatively lower performance of TF-IDF, SBERT, and BM25 in the NEI (Not Enough Information) label category. This discrepancy in performance can be attributed to the inherent challenges in dealing with the NEI label. On the NEI label, the information contained within a claim may be entirely unrelated to the available evidence. Therefore, accurately identifying and classifying such cases is inherently more complex. TF-IDF, SBERT, and BM25 struggle with this label due to the inherent ambiguity and lack of clear semantic connections between the claim and evidence.

4.2.2. Verdict Prediction

The analysis of the metrics clearly indicates that transformer-based models have emerged as the top performers in the task at hand. Their ability to capture complex linguistic patterns and relationships within the data sets them apart from other models.

In the case of the "word" corpus, PhoBERT_{Large} stands out as the best-performing model. It achieves an impressive accuracy of 82.88% on the development set and 81.63% on the test set. These results demonstrate the robustness and effectiveness of PhoBERT_{Large} in handling the nuances of the "word" corpus. Its high accuracy in both development and test sets suggests that it can consistently provide reliable predictions in this specific context.

However, when we shift our focus to the "syllable" corpus, InfoXLM_{Large} takes the lead with remarkable accuracy and F_1 scores. It achieves an accuracy of 87.45% and 86.50% on validation and test sets, respectively. These numbers not only surpass the performance of PhoBERT_{Large} but also highlight the exceptional capabilities of InfoXLM_{Large} in handling the "syllable" corpus. This outcome establishes InfoXLM_{Large} as the optimal model for the VP task. The superiority of InfoXLM_{Large} in the "syllable" corpus might be attributed to its ability to capture the intricate linguistic patterns and syllable-level features, which are evidently critical in achieving high accuracy and F_1 scores in this particular context.

	Model	Dev		Test	
		Acc	F ₁	Acc	F ₁
Word	CBoW	51.26±0.66	50.50±0.98	50.39±1.05	49.65±0.95
	BiLSTM	49.91 ±0.47	49.83 ±0.61	49.13 ±0.45	49.24 ±0.48
	ViDeBERTa _{Small}	61.67 ±1.13	61.43 ±0.79	60.62±0.66	61.73 ±0.87
	ViDeBERTa _{Base}	64.41 ±0.57	64.11 ±0.54	63.70±0.51	64.06 ±0.40
	PhoBERT _{Base}	81.72 ±0.57	81.54 ±0.58	80.67±0.37	80.70 ±0.34
	PhoBERT _{Base V2}	82.50 ±0.52	82.28 ±0.53	81.22 ±0.53	81.19 ±0.56
	PhoBERT _{Large}	82.88 ±0.97	82.74 ±1.01	81.63 ±0.47	81.62 ±0.44
Syllable	CBoW	45.25 ±0.66	44.55 ±0.63	44.24 ±1.05	43.46 ±0.64
	BiLSTM	47.68 ±0.60	47.67 ±0.77	45.12 ±0.62	45.04 ±0.79
	mBERT	75.75 ±0.68	75.61 ±1.24	75.93 ±0.58	76.01 ±1.18
	XLM-R _{Base}	79.66 ±0.95	79.48 ±0.97	78.53 ±0.41	78.56 ±0.48
	XLM-R _{Large}	86.09 ±0.72	86.02 ±0.67	85.11 ±0.27	85.15 ±0.26
	InfoXLM _{Base}	80.10 ±0.57	79.92 ±0.61	79.19 ±0.78	79.23 ±0.77
	InfoXLM _{Large}	87.45 ±0.41	87.42 ±0.60	86.50 ±0.43	86.51 ±0.62

Table 4: The result of models on Dev and Test.

ER	Top 1			Top 5		
	SUP	REF	NEI	SUP	REF	NEI
TF-IDF	61.47	58.22	34.26	83.67	81.46	54.58
SBERT	80.19	71.57	43.61	89.41	84.69	60.10
BM25	88.30	86.93	56.57	93.93	93.04	71.51

Table 5: Evidence Retrieval Results using accuracy (%).

4.2.3. Pipeline

We combine BM25 and SBERT, respectively, with the best models in the VP task, PhoBERT_{Large} and InfoXLM_{Large} for VP. The pipeline is evaluated through the test set. Table 6 gives the result of

	Pipeline	ER Acc	VP Acc	Strict Acc
SBERT	InfoXLM _{Large}	65.56	75.80	57.15
	PhoBERT _{Large}	65.56	71.54	52.84
BM25	InfoXLM _{Large}	78.38	80.63	67.00
	PhoBERT _{Large}	78.38	76.18	63.46

Table 6: Pipeline result for test set

the pipeline; the best combination for both tasks is BM25 and InfoXLM_{Large} with 67.00%, while SBERT and PhoBERT_{Large} have the lowest result. This result demonstrates the ViWikiFC corpus’s challenge for current Vietnamese fact-checking models. Besides, it provides a practical perspective on the fact-checking problem, as accurate evidence retrieval significantly influences the performance of evidence-based methods such as PhoBERT and InfoXLM. While these methods perform well in verdict prediction tasks, their strict accuracy still needs to be improved. Therefore, approaching the Vietnamese fact-checking problem using separate models for each task, first for evidence retrieval

and then for verdict prediction, needs improvement to deal with more complex claims that have been studied in English (Pan et al., 2023; Yao et al., 2023).

5. Conclusions and Future Directions

In this paper, we construct the first open-domain and high-quality Wikipedia Vietnamese corpus, including 20,916 samples to evaluate fact-checking models. Furthermore, we have discussed the data collection and annotation method and shared the insight we obtained during the annotation process, which can be applied in the development of non-English corpus creation. The best baseline and pipeline model (InfoXLM and BM25) only achieved 67.00% strict accuracy, making it a feasible challenge for the Vietnamese language model in fact-checking tasks. We believe that ViWikiFC will encourage the development of Vietnamese fact-checking research.

Following the development of fact-checking in English, we aim to develop our research for Vietnamese fact-checking by expanding our corpus from quantity to quality with more trustworthy data sources and constructing datasets not only on textual but also on images and tables for developing multi-modal fact-checking. Besides, we want our work to adapt to more advanced NLP tasks in low-resource languages, such as fake news detection or machine reading comprehension, with advanced reasoning design and sub-task functionalities.

Acknowledgement

This research is funded by Vietnam National University Ho Chi Minh City (VNU-HCM) under the grant number DS2025-26-01.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Data Availability

Data available on reasonable request.

Author Contribution

Hung Tuan Le: Conceptualization; Formal analysis; Investigation; Methodology; Validation; Visualization; Writing - review&editing. Long Truong To: Conceptualization; Data curation; Formal analysis; Investigation; Validation; Visualization; Writing - original draft. Manh Trong Nguyen: Conceptualization; Data curation; Investigation; Methodology; Writing - original draft. Kiet Van Nguyen: Conceptualization; Formal analysis; Investigation; Methodology; Validation; Supervision; Writing - review&editing.

6. Bibliographical References

- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. *Feverous: Fact extraction and verification over unstructured and structured information*. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- 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.
- Wei-Cheng Chang, X Yu Felix, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. 2019. *Pre-training tasks for embedding-based large-scale retrieval*. In *International Conference on Learning Representations*.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. *Reading Wikipedia to answer open-domain questions*. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. *Tabfact: A large-scale dataset for table-based fact verification*. In *International Conference on Learning Representations*.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. *InfoXLM: An information-theoretic framework for cross-lingual language model pre-training*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Unsupervised cross-lingual representation learning at scale*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- 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.
- Jeffrey L. Elman. 1990. *Finding structure in time*. *Cognitive Science*, 14(2):179–211.
- Joseph L. Fleiss. 1971. *Measuring nominal scale agreement among many raters*. *Psychological Bulletin*, 76:378–382.
- Alex Graves and Jürgen Schmidhuber. 2005. *Framewise phoneme classification with bidirectional lstm and other neural network architectures*. *Neural networks*, 18(5-6):602–610.
- 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-fact: 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.
- Vivek Gupta, Maitrey Mehta, Pegah Nokhiz, and Vivek Srikumar. 2020. **INFOTABS: Inference on tables as semi-structured data**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2309–2324, Online. Association for Computational Linguistics.
- Andreas Hanselowski, Christian Stab, Claudia Schulz, Zile Li, and Iryna Gurevych. 2019. **A richly annotated corpus for different tasks in automated fact-checking**. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 493–503, Hong Kong, China. Association for Computational Linguistics.
- Benjamin Horne, Sara Khedr, and Sibel Adali. 2018. Sampling the news producers: A large news and feature data set for the study of the complex media landscape. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kuebler, and Larry Moss. 2020. **Ocnli: Original chinese natural language inference**. In *Findings of EMNLP*.
- Xuming Hu, Zhijiang Guo, GuanYu Wu, Aiwei Liu, Lijie Wen, and Philip Yu. 2022. **CHEF: A pilot Chinese dataset for evidence-based fact-checking**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3362–3376, Seattle, United States. Association for Computational Linguistics.
- Tin Van Huynh, Kiet Van Nguyen, and Ngan Luu-Thuy Nguyen. 2022. **ViNLI: A Vietnamese corpus for studies on open-domain natural language inference**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3858–3872, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Yichen Jiang, Shikha Bordia, Zheng Zhong, Charles Dognin, Maneesh Singh, and Mohit Bansal. 2020. **HoVer: A dataset for many-hop fact extraction and claim verification**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3441–3460, Online. Association for Computational Linguistics.
- 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.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Jude Khouja. 2020. **Stance prediction and claim verification: An Arabic perspective**. In *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*, pages 8–17, Online. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- 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.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized BERT pre-training approach**. *CoRR*, abs/1907.11692.
- Rahmad Mahendra, Alham Fikri Aji, Samuel Louvan, Fahrurrozi Rahman, and Clara Vania. 2021. **IndoNLI: A natural language inference dataset for Indonesian**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language*

- Processing*, pages 10511–10527, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jean Maillard, Cynthia Gao, Elahe Kalbassi, Kaushik Ram Sadagopan, Vedanuj Goswami, Philipp Koehn, Angela Fan, and Francisco Guzman. 2023. [Small data, big impact: Leveraging minimal data for effective machine translation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2740–2756, Toronto, Canada. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.
- Dat Quoc Nguyen and Anh Tuan Nguyen. 2020. [PhoBERT: Pre-trained language models for Vietnamese](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1037–1042, Online. Association for Computational Linguistics.
- Linh The Nguyen and Dat Quoc Nguyen. 2021. PhoNLP: A joint multi-task learning model for Vietnamese part-of-speech tagging, named entity recognition and dependency parsing. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 1–7.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019. Multi-stage document ranking with bert. *arXiv preprint arXiv:1910.14424*.
- Jeppe Nørregaard and Leon Derczynski. 2021. [DanFEVER: claim verification dataset for Danish](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 422–428, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Liangming Pan, Xiaobao Wu, Xinyuan Lu, Anh Tuan Luu, William Yang Wang, Min-Yen Kan, and Preslav Nakov. 2023. [Fact-checking complex claims with program-guided reasoning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6981–7004, Toronto, Canada. Association for Computational Linguistics.
- Long Phan, Hieu Tran, Hieu Nguyen, and Trieu H. Trinh. 2022a. [ViT5: Pretrained text-to-text transformer for Vietnamese language generation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*, pages 136–142, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Quoc Long Phan, Tran Huu Phuoc Doan, Ngoc Hieu Le, Ngoc Bao Duy Tran, and Tuong Nguyen Huynh. 2022b. Vietnamese sentence paraphrase identification using sentencebert and phobert. In *Intelligence of Things: Technologies and Applications*, pages 416–423, Cham. Springer International Publishing.
- Chris Samarin, Wynne Hsu, and Mong Li Lee. 2020. Latent retrieval for large-scale fact-checking and question answering with nli training. In *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 941–948. IEEE Computer Society.
- Chris Samarin, Wynne Hsu, and Mong Li Lee. 2021. [Improving evidence retrieval for automated explainable fact-checking](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, pages 84–91, Online. Association for Computational Linguistics.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. [Get your vitamin C! robust fact verification with contrastive evidence](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online. Association for Computational Linguistics.
- Gautam Kishore Shahi and Durgesh Nandini. 2020. FakeCovid—a multilingual cross-domain fact check news dataset for covid-19. *arXiv preprint arXiv:2006.11343*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- 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.

Cong Dao Tran, Nhut Huy Pham, Anh Tuan Nguyen, Truong Son Hy, and Tu Vu. 2023. [ViDeBERTa: A powerful pre-trained language model for Vietnamese](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1071–1078, Dubrovnik, Croatia. Association for Computational Linguistics.

Andreas Vlachos and Sebastian Riedel. 2014. [Fact checking: Task definition and dataset construction](#). In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 18–22, Baltimore, MD, USA. Association for Computational Linguistics.

Thanh Vu, Dat Quoc Nguyen, Dai Quoc Nguyen, Mark Dras, and Mark Johnson. 2018. [Vn-CoreNLP: A Vietnamese natural language processing toolkit](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 56–60, New Orleans, Louisiana. 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.

Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nallapati, and Bing Xiang. 2019. Multi-passage bert: A globally normalized bert model for open-domain question answering. 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 5878–5882.

Barry Menglong Yao, Aditya Shah, Lichao Sun, Jin-Hee Cho, and Lifu Huang. 2023. End-to-end multimodal fact-checking and explanation generation: A challenging dataset and models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2733–2743.

A. Data Generation Rules Analysis

Rule	Example	Ratio
Change active sentences into passive sentences and vice versa	<p>E: Âu Lạc bị nhà Triệu ở phương Bắc thôn tính vào đầu thế kỷ thứ 2 TCN sau đó là thời kỳ Bắc thuộc kéo dài hơn một thiên niên kỷ. (<i>Au Lac was conquered by the Zhao in the northern region at the beginning of the 2nd century BC, leading to the Northern Domination period that lasted for over a millennium.</i>)</p> <p>C: Thời kỳ Bắc thuộc diễn ra sau khi phương Bắc thôn tính được Âu Lạc. (<i>The Northern Domination period took place after the northern regions conquered Au Lac.</i>)</p>	17.40%
Replace words with synonyms.	<p>E: Thời kỳ đầu, những bậc đế vương và những nhà quý tộc của La Mã thích lụa Trung Hoa đến mức họ cho cân lụa lên và đổi chỗ lụa đó bằng vàng với cân nặng tương đương. (<i>In the early periods, Roman emperors and nobles had such a fondness for Chinese silk that they would weigh silk against gold and exchange it with an equivalent weight of gold.</i>)</p> <p>C: Lụa Trung Hoa nhận được sự ưa chuộng cực lớn từ những người đứng đầu, những người có địa vị ở La Mã. (<i>Chinese silk received immense favor from the Roman leaders and influential people.</i>)</p>	58.60%
Add or remove modifiers that do not radically alter the meaning of the sentence.	<p>E: Gary Shilling, chủ tịch một công ty nghiên cứu kinh tế, cho rằng mức tăng trưởng GDP thực sự của Trung Quốc chỉ là 3.50% chứ không phải 7.00% như báo cáo chính thức. (<i>Gary Shilling, the chairman of an economic research firm, believes that China's actual GDP growth rate is only 3.50%, not the officially reported 7.00%.</i>)</p> <p>C: Báo cáo chính thức về mức tăng trưởng thực sự của Trung Quốc đã được Gary Shilling cho rằng là 3.50%. (<i>The official report on China's actual growth rate has been suggested by Gary Shilling to be 3.50%.</i>)</p>	48.00%
Turn nouns into relative clauses	<p>E: Trong bốn nước xưa có nền văn minh lớn thì có ba nước xưa ở vào châu Á (Ấn Độ, Iraq (Lưỡng Hà) và Trung Quốc). (<i>Among the four ancient civilizations, three were located in Asia (India, Iraq (Mesopotamia), and China).</i>)</p> <p>C: Châu Á là nơi tồn tại ba nước có nền văn minh lớn. (<i>Asia is the region where three major ancient civilizations existed.</i>)</p>	10.20%
Turn the object into relative clauses	<p>E: Lý thuyết "Quả cầu tuyết Trái Đất" cho rằng những sự thay đổi về mức độ CO2 vừa là nguyên nhân gây ra, vừa là nguyên nhân làm kết thúc thời kỳ cực lạnh ở cuối Liên đại Nguyên Sinh (Proterozoic). (<i>The "Snowball Earth" theory posits that changes in the levels of CO2 were both the cause and the termination of the extreme cold period at the end of the Proterozoic Eon.</i>)</p> <p>C: Theo lý thuyết "Quả cầu tuyết Trái Đất" cho thấy rằng sự thay đổi thành phần không khí mà cụ thể là lưu lượng CO2 là tác nhân cho kỷ băng giá, và cũng kết thúc luôn thời kỳ cực lạnh ở kỷ Proterozoic. (<i>The "Snowball Earth" theory indicates that changes in the composition of the atmosphere, specifically the levels of CO2 which are the driving factors behind ice ages and the termination of extreme cold periods in the Proterozoic Eon.</i>)</p>	2.20%

Rule	Example	Ratio
Turn adjectives into relative clauses	<p>E: Về du lịch biển, Nghệ An có 82 km bờ biển với nhều bãi tắm đẹp hấp dẫn khách du lịch quốc tế như bãi biển Cửa Lò, Cửa Hội; Nghi Thiết,... (<i>For beach tourism, Nghệ An has a coastline of 82 kilometers with numerous beautiful and attractive beaches for international tourists, such as Cua Lo, Cua Hoi, Nghi Thiet,...</i>)</p> <p>C: Nghệ An có các bãi tắm, nơi được xem là những địa điểm tuyệt đẹp và rất cuốn hút du khách quốc tế. (<i>Nghe An has beaches which are considered to be stunning and highly attractive destinations for international tourists.</i>)</p>	0.20%
Replace quantifiers or time with others that have a similar meaning.	<p>E: Theo Viện Hàn lâm Khoa học Nga, Liên Xô đã chịu 26,6 triệu thương vong trong chiến tranh thế giới thứ hai, bao gồm sự gia tăng tỷ lệ tử vong ở trẻ sơ sinh là 1,3 triệu. (<i>According to the Russian Academy of Sciences, the Soviet Union suffered 26.6 million casualties during World War II, including an increase in the infant mortality rate of 1.3 million.</i>)</p> <p>C: Liên Xô chứng kiến gần 27 triệu cái chết trong thế chiến thứ hai. (<i>The Soviet Union witnessed nearly 27 million deaths during World War II.</i>)</p>	5.60%
Create a pre-supposition sentence	<p>E: Một số người cho rằng Firenze trở thành nơi khởi đầu Phục Hưng là do may mắn, nghĩa là đơn thuần bởi vì những vĩ nhân ngẫu nhiên sinh ra ở đây: cả da Vinci, Botticelli và Michelangelo đều là người xứ Toscana (mà Firenze là thủ phủ). (<i>Some argue that Firenze became the cradle of the Renaissance due to luck, meaning it was purely coincidental that such great talents were born there: both da Vinci, Botticelli, and Michelangelo were natives of Tuscany (with Firenze as its capital).</i>)</p> <p>C: Quê hương của Michelangelo là Firenze. (<i>Michelangelo's hometown is in Firenze.</i>)</p>	20.60%
Replace Named Entities with a word that stands for the class.	<p>E: Các công ty công nghệ cao của Trung Quốc như Lenovo, Huawei, Xiaomi, Coolpad, ZTE,... đã bắt đầu cạnh tranh thành công trên thị trường thế giới. (<i>High-tech companies in China such as Lenovo, Huawei, Xiaomi, Coolpad, ZTE,... have successfully begun competing in the global market.</i>)</p> <p>C: Nhiều công ty công nghệ của Trung Quốc đã bắt đầu ghi danh mình thành công trên thị trường công nghệ cao thế giới. (<i>Many Chinese high-tech companies have indeed successfully made a name for themselves in the global high-tech market.</i>)</p>	2.00%
Create conditional sentences	<p>E: Bệnh viêm gan siêu vi C mạn được xác định là nhiễm siêu vi viêm gan C hơn 6 tháng căn cứ trên sự hiện diện của ARN. (The diagnosis of chronic hepatitis C infection is established when the presence of RNA is detected for more than 6 months.)</p> <p>C: Nếu nhiễm siêu vi gan C hơn nửa năm, bạn sẽ bị viêm gan siêu vi C mạn. (<i>If you are infected with the hepatitis C virus for more than half a year, you will be considered to have chronic hepatitis C.</i>)</p>	0.40%

Rule	Example	Ratio
Other	<p>E: Những phân tích lõi băng và lõi trầm tích đại dương không chứng minh rõ ràng sự hiện diện của băng giá và những thời kỳ trung gian băng giá trong vòng vài triệu năm qua. (<i>Core ice and sediment analyses from the ocean do not provide clear evidence of the presence of ice ages and intermediate ice-free periods over the past several million years.</i>)</p> <p>C: Dù đều là các bằng chứng địa chất nhưng theo những phân tích lõi băng và lõi trầm tích đại dương đã không chứng minh rõ ràng sự hiện diện của băng giá và những thời kỳ trung gian băng giá trong vòng vài triệu năm qua. (<i>Although both are geological pieces of evidence, core ice and ocean sediment analyses have not provided clear confirmation of the presence of ice ages and intermediate ice-free periods over the past several million years.</i>)</p> <p>Note: Although clause, Inversion can be used for converting evidence to claim. These circumstances are listed in Other.</p>	3.00%

Table 7: SUPPORTS rules and examples for creating evidence (E) - claim (C) pairs. Simply, we only mention one rule to be applied in each example.

Rule	Example	Ratio
Use negative words (no, not, never, nothing, hardly, etc.)	<p>E: Sự kiện này dẫn tới việc Hiệp định Genève (1954) được ký kết và Việt Nam bị chia cắt thành hai vùng tập kết quân sự, lấy ranh giới là vĩ tuyến 17. (This event led to the signing of the Geneva Accords in 1954, which divided Vietnam into two military zones along the 17th parallel.)</p> <p>C: Sự kiện này dẫn tới việc Hiệp định Genève (1954) không được ký kết. (<i>This event led to the Geneva Accords (1954) not being signed.</i>)</p>	23.00%
Replace words with antonyms	<p>E: Đội tuyển bóng nước Singapore đã giành huy chương vàng SEA Games lần thứ 27 vào năm 2017, tiếp tục chuỗi vô địch dài nhất của thể thao Singapore về môn môn cụ thể. (<i>The Singapore water polo team won the gold medal at the 27th SEA Games in 2017, continuing Singapore's longest winning streak in a specific sport.</i>)</p> <p>C: Singapore liên tục gặp thất bại ở bộ môn bóng nước tại các kỳ SEA Games. (<i>Singapore has continuously faced defeats in the sport of water polo at various SEA Games.</i>)</p>	14.00%
Opposite of quantity	<p>E: Năm 2010, tổng chi tiêu của Nhà nước vào khoa học và công nghệ chiếm khoảng 0.45% GDP. (<i>In 2010, the government's total expenditure on science and technology accounted for approximately 0.45% of the GDP.</i>)</p> <p>C: Nhà nước chỉ ra hơn 2.00% GDP năm 2010 cho khoa học và công nghệ. (<i>The government allocated over 2.00% of the GDP in 2010 for science and technology.</i>)</p>	9.80%
Opposite of time	<p>E: Năm 1860, dân số Singapore đã vượt quá 80,000 và hơn một nửa là người Hoa. (<i>In 1860, the population of Singapore had exceeded 80,000, with over half being of Chinese descent.</i>)</p> <p>C: Mãi đến năm 2000, Dân số Singapore mới đạt mốc 80,000. (<i>It wasn't until the year 2000 that Singapore's population reached the 80,000 milestone.</i>)</p>	5.60%

Rule	Example	Ratio
Create a sentence that has the opposite meaning of a presupposition	<p>E: Đức đã tung ra 70.00% binh lực với các sư đoàn mạnh và tinh nhuệ nhất, chưa kể binh lực góp thêm của các nước đồng minh của Đức (Ý, Rumani, Bulgari, Hungary, Phần Lan...). (<i>Germany deployed approximately 70.00% of its military power, not counting the additional military contributions from Germany's allies (Italy, Romania, Bulgaria, Hungary, Finland...).</i>)</p> <p>C: Đức không có bất kỳ đồng minh nào giúp sức. (<i>Germany did not have any allies providing assistance.</i>)</p>	5.60%
Wrong reasoning about an object (House, car, river, sea, person, etc.)	<p>E: Khi gió mùa đổi hướng, các đường bờ biển giáp với biển Ả Rập và vịnh Bengal có thể phải hứng chịu xoáy thuận. (<i>When the monsoon winds change direction, coastal areas bordering the Arabian Sea and the Bay of Bengal may be susceptible to cyclones.</i>)</p> <p>C: Các đường bờ biển tiếp giáp Ấn Độ có thể phải hứng chịu xoáy thuận khi gió mùa đổi hướng. (<i>The coastal areas bordering India may be susceptible to cyclones when the monsoon winds change direction.</i>)</p>	23.60%
Wrong reasoning about an event	<p>E: Một trận động đất ở Valdivia, Chile với cường độ 9,4-9,6 độ richter, mức cao nhất từng được ghi nhận, khiến 1.000 đến 6.000 người chết. (<i>An earthquake in Valdivia, Chile, with a magnitude of 9.4 to 9.6 on the Richter scale, the highest ever recorded, resulted in the deaths of approximately 1,000 to 6,000 people.</i>)</p> <p>C: Cơ địa chấn ở Valdivia chỉ gây ra thiệt hại về tài sản. (<i>The earthquake in Valdivia only caused property damage.</i>)</p>	30.80%
Others	<p>E: Tiếng Pháp không phải là ngôn ngữ chính thức ở Ontario, nhưng Đạo luật Dịch vụ Ngôn ngữ Pháp đảm bảo rằng các dịch vụ của tỉnh bang sẽ được cung cấp bằng ngôn ngữ này. (<i>French is not the official language in Ontario, but the French Language Services Act ensures that provincial services will be provided in this language.</i>)</p> <p>C: Dù tiếng Pháp có là ngôn ngữ chung của bang Ontario, thì nó cũng sẽ bị loại bỏ khỏi hệ thống các dịch vụ. (<i>Even though French is one of the official languages of the province of Ontario, it will be removed from the system of services.</i>)</p>	0.20%

Table 8: REFUTES rules and examples for creating evidence (E) - claim (C) pairs. , we only mention one rule to apply in each example.