

CFEVER: A Chinese Fact Extraction and VERification Dataset

Ying-Jia Lin, Chun-Yi Lin, Chia-Jen Yeh, Yi-Ting Li,
Yun-Yu Hu, Chih-Hao Hsu, Mei-Feng Lee, Hung-Yu Kao

Department of Computer Science and Information Engineering, National Cheng Kung University, Tainan, Taiwan
yingjia.lin.public@gmail.com, hykao@mail.ncku.edu.tw

Abstract

We present CFEVER, a Chinese dataset designed for **F**act Extraction and **VER**ification. CFEVER comprises 30,012 manually created claims based on content in Chinese Wikipedia. Each claim in CFEVER is labeled as “Supports”, “Refutes”, or “Not Enough Info” to depict its degree of factualness. Similar to the FEVER dataset, claims in the “Supports” and “Refutes” categories are also annotated with corresponding evidence sentences sourced from single or multiple pages in Chinese Wikipedia. Our labeled dataset holds a Fleiss’ kappa value of 0.7934 for five-way inter-annotator agreement. In addition, through the experiments with the state-of-the-art approaches developed on the FEVER dataset and a simple baseline for CFEVER, we demonstrate that our dataset is a new rigorous benchmark for factual extraction and verification, which can be further used for developing automated systems to alleviate human fact-checking efforts. CFEVER is available at <https://ikmlab.github.io/CFEVER>.

Introduction

Fact verification involves assessing the truthfulness of claims presented in text or speech. In recent years, the popularization of media platforms has accelerated the spread of misinformation, making fact verification a critical task to prevent the public from being exposed to false information. However, the process of fact verification typically involves extensive searches and assessments from many potential sources conducted by journalists, which is time-consuming and labor-intensive (Guo, Schlichtkrull, and Vlachos 2022). Consequently, it is imperative to develop automated fact verification systems to speed up the verification process.

In recent years, researchers have been building fact-verification systems with deep neural networks (Ma et al. 2016), where the models are provided with a claim and required to determine whether the claim is true or false. A well-established dataset for fact verification is FEVER (Fact Extraction and VERification) (Thorne et al. 2018a) and its shared task (Thorne et al. 2018b). FEVER asks models to extract factual sentences as evidence from the fixed Wikipedia database and provide a verdict for a given claim. Therefore, the task serves for both verification and evidence extraction for claims. Though FEVER has been widely used

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Claim: 尼克·貝爾格在 2004 年前往中東尋找生意商機。(Nick Berg traveled to the Middle East in 2004 looking for business opportunities.)

Evidence:

[[wiki/尼克·貝爾格 \(Nick Berg\)](#)]

他在 2004 年 3 月前往伊拉克尋找通訊生意商機，自 2004 年 4 月 9 日失蹤。(He went to Iraq in March 2004 to search for communication business opportunities and has been missing since April 9, 2004.)

[[wiki/伊拉克 \(Iraq\)](#)]

伊拉克共和國，通稱伊拉克，位於西亞—中東地區的共和國。(The Republic of Iraq, commonly known as Iraq, is a republic located in West Asia - the Middle East.)

Verdict: Supports

Table 1: An example from the CFEVER dataset. Underlined words indicate knowledge that should be verified from other pages in Wikipedia.

as a benchmark for fact verification, the dataset and most of the other fact verification datasets (Wang 2017; Hanselowski et al. 2019; Schuster, Fisch, and Barzilay 2021) are created in English. The dissemination of rumors and fake news is a serious problem in East Asia, especially in China. Compared to English, text written in Chinese tends to be more ambiguous and nuanced, making it more difficult for people to identify misinformation. Therefore, it is critical to create a clean and high-quality dataset as supervised data for Chinese fact verification. To achieve this goal, Hu et al. (2022) proposed the CHEF dataset from the sources of the fact-checking and news websites. However, compared to English fact verification datasets (Augenstein et al. 2019; Kotonya and Toni 2020), the size of CHEF is relatively small, making training and evaluating fact verification systems challenging. A potential alternative for building Chinese fact verification systems is to take advantage of the multilingual fact-checking dataset (Nielsen and McConville 2022). Nevertheless, the Chinese claims in the dataset are still limited in size.

Dataset	#Claims	Source	Language	Task		
				Document Retrieval	Evidence Retrieval	Claim Verification
Weibo-16 (Ma et al. 2016)	4,664	Weibo	Chinese	No	No	Yes*
Weibo-20 (Zhang et al. 2021)	6,362	Weibo	Chinese	No	No	Yes*
MuMiN (Nielsen and McConville 2022)	1,283 [†]	Tweets	Multi [‡]	No	No	Yes
CHEF (Hu et al. 2022)	10,000	Fact-checking/ News websites	Chinese	No	Yes	Yes
CFEVER	30,012	Wikipedia	Chinese	Yes	Yes	Yes

* Rumor detection task. [†] Including Chinese and other languages. [‡] Multi-lingual dataset.

Table 2: Comparison of CFEVER with other fact verification datasets in Chinese.

In this work, we present CFEVER, a Chinese dataset for **F**actual **E**xtraction and **VER**ification. Following FEVER’s construction process (Thorne et al. 2018a), we first built a fact database with the fixed version of the Chinese Wikipedia dump. Then, we hired several workers from our university to alter the extracted sentences from the Wikipedia pages and label each claim with “Supports”, “Refutes”, or “Not Enough Info” based on the Wikipedia pages in the fact database. The relevant factual sentences from the fact database were also annotated by our workers as evidence for the claims belonging to the first two categories. Our CFEVER dataset includes 30,012 claims in total, which are three times the size of CHEF (Hu et al. 2022). To evaluate our labeling quality, we report the five-way inter-annotator agreement of 0.7934 in Fleiss κ (Fleiss 1971) using 6.45% claims of our dataset, while the scores in FEVER and CHEF are 0.6841 and 0.74, measured from 4% and 3% of claims from the datasets, respectively.

To thoroughly assess the challenges of CFEVER, following the FEVER task (Thorne et al. 2018b), we conduct experiments on the three stages: document retrieval, sentence retrieval, and claim verification. We test the performance for each stage and the full-pipeline setting with the state-of-the-art approaches (Stammach 2021; DeHaven and Scott 2023) developed on FEVER and our simple baseline (Robertson, Zaragoza et al. 2009; Soleimani, Monz, and Worring 2020) along with oracle experiments to dive into the challenges of CFEVER. Further extensive analysis even reveals the characteristics and difficulty of CFEVER. In summary, we list our contributions as follows:

- We present CFEVER, the currently largest Chinese dataset for evidence-based fact verification.
- The five-way inter-annotator agreement of CFEVER in Fleiss κ indicates that the dataset was built with high label consistency among claims.
- Extensive experiments on CFEVER show that the proposed dataset can serve as a challenging benchmark for future research or development on Chinese fact verification.

Related Work

FEVER

Thorne et al. (2018b) created the task of Fact Extraction and

VERification (FEVER) along with the dataset of the same name (Thorne et al. 2018a). FEVER is the currently largest fact verification dataset that contains 185,445 claims. Each of the claims was first sampled from the sentences in the introductory sections of approximately 50,000 popular pages and then revised by human annotators. In another round of annotation, the annotators label the claims as “Supports”, “Refutes”, or “Not Enough Information”, and discover evidence sentences from Wikipedia pages. The following year, Thorne et al. (2019) introduced the FEVER 2.0 task to first generate adversarial claims to fool the existing verification systems built with the FEVER dataset (Thorne et al. 2018a). Then the task required participants to improve the systems to prevent such adversarial attacks. More recently, Aly et al. (2021) proposed FEVEROUS and extended the fact verification task for verifying claims with the information in structured data, such as tables in Wikipedia.

Nørregaard and Derczynski (2021) follow the annotation process of FEVER and propose a new FEVER dataset (Thorne et al. 2018a) in Danish. Jiang et al. (2020) report that 87% of the claims in the FEVER dataset (Thorne et al. 2018a) require only a single Wikipedia page for verification, which does not support real-world situations where misinformation may come from multiple articles. Thus, Jiang et al. (2020) propose a new dataset containing 26K claims for multi-hop reasoning based on the building process of FEVER.

Chinese Fact Verification

We compare CFEVER with other Chinese fact verification datasets in Table 2. Existing Chinese fact verification datasets mainly focus on rumor detection, such as the Weibo-16 (Ma et al. 2016) and Weibo-20 datasets (Zhang et al. 2021). In this work, we focus on both the fact extraction and verification tasks, which are different from the binary claim detection in rumor detection (Guo, Schlichtkrull, and Vlachos 2022). The dataset closest to our work is CHEF (Hu et al. 2022), which is a pilot Chinese dataset for evidence-based fact-checking. There are two main differences between CHEF and our dataset. First, CHEF is created from the sources of fact-checking websites, while our dataset is created from Wikipedia. Second, there is no document retrieval process in the task of CHEF, with candidate evidence sentences provided for each claim in the dataset. In contrast,

we follow FEVER (Thorne et al. 2018a) to provide a fixed fact database and ask models to first extract relevant documents from Wikipedia before verifying the claims with the evidence sentences. We consider our approach, along with FEVER (Thorne et al. 2018a), to be more realistic for real-world fact verification scenarios.

Dataset Construction

We followed the labeling approach of FEVER (Thorne et al. 2018a) to create the CFEVER dataset and adapted the annotation platform¹ publicly released by Thorne et al. (2018a) for our construction task. The annotation process consists of two stages: **claim generation** and **claim annotation**. Both stages are conducted by human workers recruited from our university. The two stages of annotation are conducted separately. To distinguish the two stages, we refer to the workers in the claim generation stage as **writers** and the ones in the claim annotation stage as **annotators**, based on the characteristics of their tasks. Before the construction process, we first describe our method to prepare the Wikipedia data and the fact database.

Preparation for the Wikipedia Data

We used the December 2022 dump of the Chinese Wikipedia and extracted the text from the introductory section from each page following the pre-processing method of FEVER (Thorne et al. 2018a). The raw Wikipedia comprises articles in both Traditional and Simplified Chinese. To unify the data, we processed the text with the open-source software OpenCC² to convert the text to Traditional Chinese. Then, the processed data containing 1,187,751 pages were fixed to serve as the fact database for CFEVER. For the next stage of data construction, we created a source page pool based on the processed Wikipedia pages. The source page pool includes the 500 most visited Chinese Wikipedia pages³ worldwide in 2022, 10,000 Taiwanese pages, and 3,000 random pages. All the claims in our dataset were created based on the pages in the source page pool.

Claim Generation

At this stage, writers are responsible for writing claims based on the Wikipedia pages in the source page pool. At each time, each writer was given one extracted sentence randomly sampled from the introductory section of a page in the source page pool. In addition, Wikipedia pages related to the given sentence based on the hyperlinks in the raw Wiki data were also provided to the writer. Then, we asked a writer to first generate a TRUE claim based on the given extracted sentence and the information in the relevant pages without any learned human knowledge. The provided relevant pages were used to help writers come up with diverse claims, which may also result in complex claims that require multi-step reasoning from different pages in the verification

¹<https://github.com/aws-labs/fever/tree/master/fever-annotations-platform>

²<https://github.com/BYVoid/OpenCC>

³<https://pageviews.wmcloud.org/topviews/?project=zh.wikipedia.org>




Figure 1: Domain distribution of the generated claims in our dataset. H&S refers to the domain of Humanities & Social Sciences.

task. After that, the writer was asked to generate six variants of the TRUE claim:

- **Rephrasing:** A TRUE claim should be rephrased to a different sentence with the same meaning.
- **Negation:** A TRUE claim should be negated without simple negation words, such as “not”.
- **Entity substitution at a similar level:** An entity in a TRUE claim should be substituted with another one similar to the original entity.
- **Entity substitution at a disjointed level:** An entity in a TRUE claim should be substituted with another one disjointed to the original entity.
- **Specification:** A TRUE claim should be narrowed down with more specific concepts.
- **Generalization:** A TRUE claim should be generalized with more abstract concepts.

During claim generation, the writers were asked not to generate claims with their own learned knowledge. They should write claims solely based on the information in the given extracted sentence and the relevant pages. The reason behind this step is to help generate verifiable claims and maintain the quality of the generated claims among different writers. We also measured the domains of the claims in CFEVER. To prevent misclassification based on the domains of source Wiki pages, we asked the writers to select a category from the pre-defined domains for each claim they generated. The final domain distribution of the generated claims is shown in Figure 1.

Claim Annotation

Once a claim was generated, an annotator was asked to label the claim as “Supports”, “Refutes”, or “Not Enough Info”. For the first two categories, the annotator must also find the sentences as evidence from the fact database. To achieve such a process, at each time of annotation, four kinds of materials are provided to the annotator in default:

- **Claim:** The claim generated by the writer in the claim generation stage.
- **Page name:** The title of the original page from which the claim was generated.
- **Original sentences:** The sentences in the introductory section of the original page which the claim was generated from.
- **Relevant pages:** The pages related to the original page based on the hyperlinks in the raw Wiki data.

If none of the sentences provided in default can be selected as evidence for the claim, an annotator can also search the fact database with a Wiki page name as a keyword. After passing the keyword to the annotation platform, the sentences in the introductory section of the Wikipedia page will show up at the annotation interface to be selectable as evidence by an annotator. Note that a claim may become “Supports” or “Refutes” based on multiple sentences. We encouraged the annotators to find as many sentences as possible to support or refute a claim from the fact database. Once the annotator considered no factual sentences in our fact database could support or refute a claim, the claim was labeled as “Not Enough Info.”

Some claims were generated based on the original extracted sentence along with the content from the relevant pages in the claim generation stage. These claims require evidence referenced from multiple pages. For the example in Table 1, the term “Middle East” in the claim cannot be inferred directly from the sentence mentioning “Iraq.” Another sentence in the “Iraq” page should be selected as a part of the evidence, even though the relationship between the two terms is common knowledge. Therefore, for these claims, we asked annotators to combine two or more sentences from different pages as complete evidence to support or refute the claim using the annotation platform.

Workers

We recruited nine writers and annotators from our university in total. All of the workers were native Chinese speakers. Among them, three workers were from the College of Liberal Arts, with two of them focused on the claim generation task only, and the remaining six were from the College of Engineering. The workers were trained by the authors for the annotation tasks and guidelines until they were able to undergo the annotation process correctly and independently.

Data Validation

To evaluate the consistency of the class labels among the annotators, we randomly sampled 1,936 claims (6.45% of our total dataset) and asked five annotators from our workers to label them. The five-way inter-annotator agreement over the 1,936 claims shows a score of 0.7934 in Fleiss κ (Fleiss 1971), which is higher than the score of 0.6841 measured with 7,506 claims (4% from total) in FEVER (Thorne et al. 2018a) and the one of 0.74 with 310 claims (3% from total) in the CHEF dataset (Hu et al. 2022). To evaluate the correctness of the evidence sentences, we randomly sampled another 700 claims (2.33% of the dataset) for being reviewed

Split	Training	Development	Test
Total Claims	24,012	3,000	3,000
Num of SUP	11,085	1,000	1,000
Num of REF	7,113	1,000	1,000
Num of NEI	5,814	1,000	1,000
Avg. Claim Length	33.19	34.04	34.04
Avg. Evidence per Claim	1.57 sents	1.53 sents	1.55 sents
Avg. Evidence Pages per Claim	1.13 (88.41%)	1.14 (87.35%)	1.12 (89.45%)

Table 3: Dataset statistics of CFEVER in different splits. SUP indicates the “Supports” class, REF stands for “Refutes”, and NEI represents “Not Enough Information”. The Avg. Claim Length is the average number of characters in a claim. The Avg. Evidence per Claim is the average number of evidence sentences for a claim in the “Supports” or “Refutes” class, and the ratio in the parentheses is the proportion of claims with evidence from one single page.

by the authors. We discovered that 84.4% of the claims are annotated correctly with correct evidence sentences.

Dataset Statistics

There are 30,012 claims in the CFEVER dataset. We split 80%, 10%, and 10% of the claims into the training, development, and test sets, respectively. The statistics of the dataset are shown in Table 3. The number of claims in the three categories is balanced in the development and test sets to ensure that the performance of the models is not biased towards any category during evaluations. The average claim length (character level) and the average number of evidence sentences per claim are also similar among the three splits. We also report the ratio of claims with evidence from a single page in Table 3, with 88.41%, 87.35%, and 89.45% of the claims whose evidence can be found in a single page in the three sets. These ratios are close to the statistic of FEVER (Thorne et al. 2018a) reported by Jiang et al. (2020), where 87% of the claims require one page for verification.

Baseline Systems

Following Thorne et al. (2018b), our task requires a model to retrieve evidence from the Wikipedia fact database and perform verification for each claim. This section introduces the approaches we test for CFEVER in three stages: document retrieval, sentence retrieval, and recognizing textual entailment (RTE) for claim verification. The systems we build involves two full-pipeline methods⁴ for the three stages: one simple baseline proposed by ourselves and the state-of-the-art approach (DeHaven and Scott 2023) developed for FEVER (Thorne et al. 2018a). We also test CFEVER with the sentence retrieval approach proposed by Stammbach (2021). We provide essential details in this section, and more information is available on the CFEVER website.

⁴BERT (<https://huggingface.co/hfl/chinese-bert-wwm-ext>) is used for BEVERS and our baseline unless otherwise noted.

Our Baseline

To understand the difficulty and behaviors of CFEVER, we first design a baseline with simple components to explore the dataset.

Document Retrieval For evidence-based claim verification, relevant pages should be discovered for each claim to extract evidence. Many studies for the FEVER task (Thorne et al. 2018b) adopt TF-IDF (Thorne et al. 2018a) or the search with the MediaWiki API (Hanselowski et al. 2018) for document retrieval. Following Jiang et al. (2020), we use BM25 (Robertson, Zaragoza et al. 2009) for retrieving relevant pages for each claim. Our implementation was based on Elasticsearch⁵, and the representations were built with the Wikipedia pages from our fact database.

Sentence Retrieval After retrieving relevant pages for each claim, we perform sentence retrieval to select evidence sentences from the pages. Inspired by Hanselowski et al. (2018) and Soleimani, Monz, and Worring (2020), we implement a pointwise approach for sentence retrieval with BERT (Devlin et al. 2019) to classify each claim-sentence pair in binary. Positive pairs are created using claims and their corresponding gold evidence sentences. In contrast, negative pairs consist of claims paired with non-evidence sentences, which are sampled from the predicted pages acquired during the document retrieval phase.

Recognizing Textual Entailment We verify each claim with the evidence sentences by recognizing textual entailment (RTE). Following Hanselowski et al. (2018); Nie, Chen, and Bansal (2019), we first concatenate a claim with its top five evidence sentences and then fine-tune the BERT model (Devlin et al. 2019) for the three-class RTE task.

BEVERS

Document Retrieval The second approach is based on the BEVERS (DeHaven and Scott 2023), which is the state-of-the-art full-pipeline system for FEVER (Thorne et al. 2018b). BEVERS uses a hybrid approach to include both the search results from Wikisearch (Hanselowski et al. 2018) and the TF-IDF method (Thorne et al. 2018a). Such a hybrid approach was also adopted by Stammbach (2021). Note that BEVERS replaced the MediaWiki API in the approach of Hanselowski et al. (2018) with a fuzzy string search system.

Sentence Retrieval Besides the document retrieval approach, DeHaven and Scott (2023) also proposed a competitive sentence retrieval approach. BEVERS extends the binary pointwise approach (Hanselowski et al. 2018) with an additional ternary classification task for classifying a claim-sentence pair into “Supports”, “Refutes”, or “Not Enough Info”, as an initial set of predicted evidence sentences. Then, BEVERS used the results from the initial set to explore more evidence sentences from the hyperlinks in a Wikipedia article, which was called “re-retrieval.” Finally, all extracted sentences from these two steps are ranked to yield the final evidence sentences for each claim.

⁵<https://github.com/elastic/elasticsearch>

Recognizing Textual Entailment In addition to the concatenation-based approach (Hanselowski et al. 2018; Nie, Chen, and Bansal 2019) for performing claim verification with concatenated evidence, the singleton-based approach (Malon 2018; Soleimani, Monz, and Worring 2020) was also proposed to classify each claim-evidence pair individually. In this setting, each claim will have multiple scores for each evidence sentence. Then, the scores are aggregated based on the rules (Malon 2018; Soleimani, Monz, and Worring 2020) to obtain the final prediction. BEVERS adopts a mixture of both approaches (DeHaven and Scott 2023). They first fine-tune DeBERTa-V2-XL (He et al. 2021) pre-trained on MNLI (Williams, Nangia, and Bowman 2018) for each of the approaches and train an additional gradient boosting classifier (Friedman 2001) for aggregating the final predictions.

Stammbach

Stammbach (2021) treats sentence retrieval as a token-level classification problem, where a model must predict 1 for each token within evidence sentences and 0 for the tokens belonging to non-evidence sentences. Such an approach requires a model to process long sequences from input claim-article pairs. Thus, Stammbach (2021) adopts BigBird (Zaheer et al. 2020) as the encoder. Since this approach was proposed for sentence retrieval, we only test this approach for retrieving evidence with the ground-truth documents.

Evaluation Metrics

For document retrieval and sentence retrieval, we report the performance in recall (%). Our recall evaluation metric is designed to assess the model’s ability to correctly predict at least one complete set of evidence pages during document retrieval and, similarly, at least one complete set of evidence sentences during sentence retrieval, for each data instance. For claim verification in RTE, following Thorne et al. (2018b), we report performance in accuracy (%) and FEVER Score (%). The latter is a strict measure of accuracy, requiring a model to correctly predict at least one complete evidence set for each claim. For implementing the evaluation metrics, we use the script from DeHaven and Scott (2023) for document retrieval. For sentence retrieval and RTE, we utilize the official scoring tool from Thorne et al. (2018b).

Results and Analysis

Results in Different Stages

To thoroughly analyze CFEVER, we list the results for the different stages in Table 4. For document retrieval, we observe that BEVERS achieves more than 90% of recall, demonstrating that the approach is able to find correct articles for most of the given claims. The performance difference between our simple baseline (BM25) and BEVERS is primarily due to the hybrid approach employed in BEVERS, which combines the search predictions from the TF-IDF method and the fuzzy string search.

For sentence retrieval, there’s a huge performance gap between our simple baseline and BEVERS. This may result from the training with an additional ternary classifica-

tion task and the employment of the re-retrieval technique in BEVERS, whose score is also far ahead of the other baselines close to our sentence selection method on FEVER. However, BEVERS scored 94.41% in recall reported in their paper (DeHaven and Scott 2023) for FEVER, showing that our dataset remains a challenge for BEVERS.

For claim verification and the full-pipeline setting, BEVERS obtains only 69.73% for the label accuracy, which is much lower than their score (80.24%) reported for the FEVER dataset (DeHaven and Scott 2023). Since the classification is based on the evidence extracted in the sentence retrieval stage, the scores for claim verification will be affected if a model cannot identify correct evidence sentences. Still, BEVERS significantly outperforms our simple baseline by 8.56% in label accuracy and 12.33% in FEVER Score.

We also test the performance of GPT-3.5 (Ouyang et al. 2022) and GPT-4 (OpenAI 2023)⁶ for claim verification with the test set using the *zero-shot* and few-shot settings. For the few-shot setting, we sample three labeled claims from the training set with the same domain as the input claim for each class. From Table 4, we find that the claims in CFEVER are challenging for both models, and the performance can be slightly improved with the few-shot setting.

Oracle Results

To analyze the difficulty of CFEVER further, we also report the oracle results for the last two stages in Table 5. For the oracle setting in sentence retrieval, the gold documents of each claim are provided for the models. As for the oracle setting in claim verification, the models take gold evidence sentences for each claim as inputs. The results show that both methods achieve more than 95% in recall for the oracle setting in sentence retrieval. However, there are still 15% and 10% error rates in claim verification for our baseline and BEVERS (RoBERTa_{Large}⁷), showing that certain claims in our dataset pose verification challenges. We notice that Stammbach performs worse than the other two baselines for sentence retrieval. The results may be affected by the Chinese BigBird we used⁸, since there is no official Chinese version of BigBird (Zaheer et al. 2020).

Analysis for Claims with Evidence from Multiple Pages

After testing the two methods in different stages, we further analyze the performance on claims with evidence labeled from different numbers of Wikipedia pages for the full-pipeline setting in accuracy and FEVER Score. The results in Table 6 show that the performance of both methods degrades significantly for the claims with evidence from more pages. For the claims with more than three pages of evidence, the two methods achieve unsatisfactory performance in FEVER Score. In summary, we identify that about 10% claims with evidence from multiple pages in our dataset are challenging to verify. Even BEVERS can only get about

⁶The prompts we used can be found from the CFEVER website.

⁷<https://huggingface.co/hfl/chinese-robetta-wwm-ext-large>

⁸<https://huggingface.co/Lowin/chinese-bigbird-base-4096>

Task (metric)	System	Score (%)
Doc retrieval (Recall)	Our baseline	87.65
	BEVERS ^a	92.60
Sent retrieval (Recall)	Our baseline	76.65
	BEVERS ^a	86.60
RTE (Accuracy)	Our baseline	61.17
	BEVERS ^a	69.73
	GPT-3.5 (<i>zeroshot</i>)	43.17
	GPT-3.5 (<i>3-shot</i>)	44.20
	GPT-4 (<i>zeroshot</i>)	47.23
Full pipeline (FEVER Score)	Our baseline	52.47
	BEVERS ^a	64.80

Table 4: Results on the different stages. *a*: DeHaven and Scott (2023).

Task (metric)	System	Score (%)
Sent retrieval (Recall)	Our baseline	95.90
	BEVERS ^a	95.20
	Stammbach ^b	83.55
RTE (Accuracy)	Our baseline	85.10
	BEVERS ^a (BERT _{Base})	88.50
	BEVERS ^a (RoBERTa _{Large})	90.33

Table 5: Oracle results on sentence retrieval and the task of recognizing textual entailment (RTE). *a*: DeHaven and Scott (2023). *b*: Stammbach (2021).

70% of the FEVER Score for the claims with evidence from single pages.

Analysis for Claims with Different Numbers of Evidence Sentences

Since each claim in our dataset can have multiple evidence sentences, we also analyze the performance on claims with different numbers of evidence sentences. The results in Table 7 show that both methods perform better on the claims with two and three evidence sentences in the full-pipeline setting. This may result from the fact that the claims with one evidence sentence are usually shorter ones, containing less information to be verified. Additionally, the claims with more than three evidence sentences are usually much longer, requiring more correct evidence sentences to be identified and thus more difficult for the models.

Analysis for Claims from Different Domains

As the chart in Figure 1 shows, the claims in our dataset were labeled into 11 domains by our writers during the claim generation process. Figure 2 shows the performance of the two methods in FEVER Score on claims from different domains. We discover that both of the methods have lower performance for the domains with fewer training claims, such as

Methods	#Pages (ratio)		
	1 (89.45%)	2 (9.35%)	≥ 3 (1.20%)
Ours	67.75 / 57.97	55.08 / 15.51	50.00 / 0.00
BEVERS ^a	73.95 / 70.54	63.10 / 21.93	58.33 / 16.67

Table 6: Full-pipeline results in accuracy / FEVER Score (%) for the claims with evidence from different number of pages. “Ratio” indicates the proportion of the claims from the test set (not counted for the claims of “Not Enough Information”) in each group. *a*: DeHaven and Scott (2023).

Methods	#Evidence sentences (ratio)			
	1 (61.75%)	2 (26.40%)	3 (10.15%)	≥ 4 (1.70%)
Ours	64.1 / 55.9	72.2 / 56.3	65.5 / 36.9	61.8 / 11.8
BEVERS ^a	70.8 / 67.8	75.6 / 65.3	80.3 / 57.1	55.9 / 26.5

Table 7: Full-pipeline results in accuracy / FEVER Score (%) for the claims with different numbers of evidence sentences. “Ratio” indicates the proportion of the claims from the test set (not counted for the claims of “Not Enough Information”) in each group. *a*: DeHaven and Scott (2023).

“Sports,” “Technology,” and “Politics.” We also observe that the number of Wikipedia pages of the evidence for claims in the “Sports” domain is higher than the numbers of all the other domains, which may result in lower performance and match our findings in Table 6.

Analysis for Claims in Different Lengths

In this section, we analyze the performance of the two methods on claims with different lengths in the full-pipeline setting. We divide the claims in the test set into five groups according to the number of characters in the claims and show the results in Figure 3. We observe that our simple baseline performs worse on the claims longer than 51 characters, while BEVERS remains stable. Furthermore, the two methods perform better on the medium length claims with 31-40 characters. These results again indicate that the length of the claims is an important factor in the claim verification task.

Discussion

We discuss two limitations of our dataset in this section. First, although CFEVER is currently the largest Chinese dataset for fact extraction and verification, it is still much smaller than the English FEVER dataset. Data size is an important factor for model performance and generalization. We hope to increase the data scale in the future while maintaining the high quality of annotation. Second, CFEVER may not be a perfect dataset for training models to handle complex reasoning tasks. As we show in Table 2, most of the claims in CFEVER require evidence extracted from only one Wikipedia page. This issue has also been reported for FEVER (Jiang et al. 2020). We hope to expand our dataset with more complex claims that require evidence from multiple pages in the future.

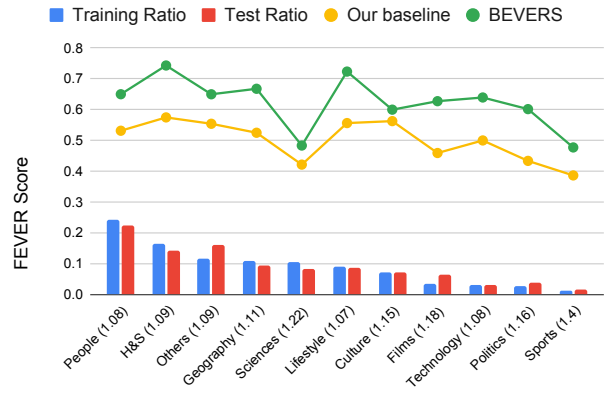


Figure 2: Performance comparisons for the claims of different domains in the full-pipeline setting. H&S refers to the domain of Humanities & Social Sciences. Values in the parentheses are the average number of evidence pages.

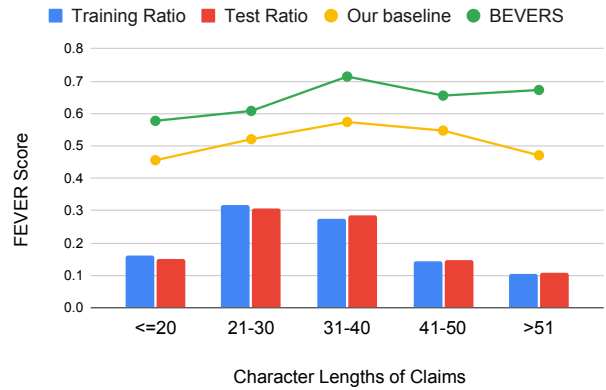


Figure 3: Performance comparisons with different lengths of claims in the full-pipeline setting. “Ratio” indicates the proportion for each group of the claims from the test set.

Conclusion

This paper introduces CFEVER, a new Chinese dataset for Fact Extraction and VERification. Following the FEVER task (Thorne et al. 2018b), CFEVER forms a verification task that requires models to verify the claims into “Supports”, “Refutes”, and “Not Enough Information”. In addition, for the first two categories, models are also required to extract the evidence sentences from our fact database composed of Chinese Wikipedia pages. We carefully validate the quality of the dataset and obtain an inter-annotator agreement of 0.7934 in Fleiss κ for the class label consistency. Though the experiments with the simple baseline designed by ourselves and the state-of-the-art method (DeHaven and Scott 2023) developed on FEVER (Thorne et al. 2018a), we believe that CFEVER is a challenging dataset to serve as a benchmark on Chinese claim verification and fact extraction.

Acknowledgements

This work was supported by the National Science and Technology Council of Taiwan, under Grant NSTC 112-2223-E-006-009. We thank the anonymous reviewers for their insightful comments. We also extend our deepest gratitude to all of our data annotators for their hard work and dedication during the data construction process.

References

- Aly, R.; Guo, Z.; Schlichtkrull, M. S.; Thorne, J.; Vlachos, A.; Christodoulopoulos, C.; Cocarascu, O.; and Mittal, A. 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)*.
- Augenstein, I.; Lioma, C.; Wang, D.; Chaves Lima, L.; Hansen, C.; Hansen, C.; and Simonsen, J. G. 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)*, 4685–4697. Hong Kong, China: Association for Computational Linguistics.
- DeHaven, M.; and Scott, S. 2023. BEVERS: A General, Simple, and Performant Framework for Automatic Fact Verification. In *Proceedings of the Sixth Fact Extraction and VERification Workshop (FEVER)*, 58–65. Dubrovnik, Croatia: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Fleiss, J. L. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5): 378.
- Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 1189–1232.
- Guo, Z.; Schlichtkrull, M.; and Vlachos, A. 2022. A Survey on Automated Fact-Checking. *Transactions of the Association for Computational Linguistics*, 10: 178–206.
- Hanselowski, A.; Stab, C.; Schulz, C.; Li, Z.; and Gurevych, I. 2019. A Richly Annotated Corpus for Different Tasks in Automated Fact-Checking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, 493–503. Hong Kong, China: Association for Computational Linguistics.
- Hanselowski, A.; Zhang, H.; Li, Z.; Sorokin, D.; Schiller, B.; Schulz, C.; and Gurevych, I. 2018. UKP-Athene: Multi-Sentence Textual Entailment for Claim Verification. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, 103–108. Brussels, Belgium: Association for Computational Linguistics.
- He, P.; Liu, X.; Gao, J.; and Chen, W. 2021. DEBERTA: DECODING-ENHANCED BERT WITH DISENTANGLED ATTENTION. In *International Conference on Learning Representations*.
- Hu, X.; Guo, Z.; Wu, G.; Liu, A.; Wen, L.; and Yu, P. 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*, 3362–3376. Seattle, United States: Association for Computational Linguistics.
- Jiang, Y.; Bordia, S.; Zhong, Z.; Dognin, C.; Singh, M.; and Bansal, M. 2020. HoVer: A Dataset for Many-Hop Fact Extraction And Claim Verification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 3441–3460. Online: Association for Computational Linguistics.
- Kotonya, N.; and Toni, F. 2020. Explainable Automated Fact-Checking for Public Health Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 7740–7754. Online: Association for Computational Linguistics.
- Ma, J.; Gao, W.; Mitra, P.; Kwon, S.; Jansen, B. J.; Wong, K.-F.; and Cha, M. 2016. Detecting Rumors from Microblogs with Recurrent Neural Networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, 3818–3824. AAAI Press. ISBN 9781577357704.
- Malon, C. 2018. Team Papelo: Transformer Networks at FEVER. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, 109–113. Brussels, Belgium: Association for Computational Linguistics.
- Nie, Y.; Chen, H.; and Bansal, M. 2019. Combining fact extraction and verification with neural semantic matching networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 6859–6866.
- Nielsen, D. S.; and McConville, R. 2022. MuMiN: A Large-Scale Multilingual Multimodal Fact-Checked Misinformation Social Network Dataset. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, 3141–3153. New York, NY, USA: Association for Computing Machinery. ISBN 9781450387323.
- Nørregaard, J.; and Derczynski, L. 2021. DanFEVER: claim verification dataset for Danish. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, 422–428. Reykjavik, Iceland (Online): Linköping University Electronic Press, Sweden.
- OpenAI, R. 2023. GPT-4 technical report. *arXiv*, 2303–08774.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds.,

- Advances in Neural Information Processing Systems*, volume 35, 27730–27744. Curran Associates, Inc.
- Robertson, S.; Zaragoza, H.; et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4): 333–389.
- Schuster, T.; Fisch, A.; and Barzilay, R. 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*, 624–643. Online: Association for Computational Linguistics.
- Soleimani, A.; Monz, C.; and Worring, M. 2020. BERT for Evidence Retrieval and Claim Verification. In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II*, 359–366. Berlin, Heidelberg: Springer-Verlag. ISBN 978-3-030-45441-8.
- Stammbach, D. 2021. Evidence Selection as a Token-Level Prediction Task. In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, 14–20. Dominican Republic: Association for Computational Linguistics.
- Thorne, J.; Vlachos, A.; Christodoulopoulos, C.; and Mittal, A. 2018a. 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)*, 809–819. New Orleans, Louisiana: Association for Computational Linguistics.
- Thorne, J.; Vlachos, A.; Cocarascu, O.; Christodoulopoulos, C.; and Mittal, A. 2018b. The Fact Extraction and VERification (FEVER) Shared Task. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, 1–9. Brussels, Belgium: Association for Computational Linguistics.
- Thorne, J.; Vlachos, A.; Cocarascu, O.; Christodoulopoulos, C.; and Mittal, A. 2019. The FEVER2.0 Shared Task. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, 1–6. Hong Kong, China: Association for Computational Linguistics.
- Wang, W. Y. 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)*, 422–426. Vancouver, Canada: Association for Computational Linguistics.
- Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1112–1122. New Orleans, Louisiana: Association for Computational Linguistics.
- Zaheer, M.; Guruganesh, G.; Dubey, K. A.; Ainslie, J.; Alberti, C.; Ontanon, S.; Pham, P.; Ravula, A.; Wang, Q.; Yang, L.; and Ahmed, A. 2020. Big Bird: Transformers for Longer Sequences. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 17283–17297. Curran Associates, Inc.
- Zhang, X.; Cao, J.; Li, X.; Sheng, Q.; Zhong, L.; and Shu, K. 2021. Mining Dual Emotion for Fake News Detection. In *Proceedings of the Web Conference 2021, WWW '21*, 3465–3476. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383127.