CHEF: A Pilot Chinese Dataset for Evidence-Based Fact-Checking

Xuming Hu[1*], Zhijiang Guo[2*], Guanyu Wu[1], Aiwei Liu[1], Lijie Wen[1†], Philip S. Yu[1,3]
[1]Tsinghua University
[2]University of Cambridge
[3]University of Illinois at Chicago
[1]{hxm19,wugy18,liuaw20}@mails.tsinghua.edu.cn
[2]zg283@cam.ac.uk [1]wenlj@tsinghua.edu.cn [3]psyu@uic.edu

Abstract

The explosion of misinformation spreading in the media ecosystem urges for automated fact-checking. While misinformation spans both geographic and linguistic boundaries, most work in the field has focused on English. Datasets and tools available in other languages, such as Chinese, are limited. In order to bridge this gap, we construct CHEF, the first CHinese Evidence-based Fact-checking dataset of 10K real-world claims. The dataset covers multiple domains, ranging from politics to public health, and provides annotated evidence retrieved from the Internet. Further, we develop established baselines and a novel approach that is able to model the evidence retrieval as a latent variable, allowing jointly training with the veracity prediction model in an end-to-end fashion. Extensive experiments show that CHEF will provide a challenging testbed for the development of fact-checking systems designed to retrieve and reason over non-English claims. Source code and data are available[1].

1 Introduction

Misinformation is being spread online at increasing rates, posing a challenge to media platforms from newswire to social media. In order to combat the proliferation of misinformation, fact-checking is an essential task that assesses the veracity of a given claim based on evidence (Vlachos and Riedel, 2014). Fact-checking is commonly conducted by journalists. However, fact-checking is a time-consuming task, which can take journalists several hours or days (Adair et al., 2017). Thus, there is a need for automating the process.

Although misinformation spans both geographic and linguistic boundaries, most existing works focused on English (Wang, 2017; Thorne et al., 2018; Augenstein et al., 2019; Hanselowski et al., 2019;

Claim: 2019年，共有12.08万人参加成都中考，但招生计划只有4.3万。 *In 2019, a total of 120,800 students participated in the high school entrance examination in Chengdu, but schools only enrolled 43,000 students.*

Document: 今年共有12.08万人参加中考，这个是成都全市，包括了20个区，高新区和天府新区的总参考人数。月前，教育局公布了2019年的普高招生计划。招生计划数进一步增加，上普高的机会更大了... 中心城区（13个区）招生计划为43015人。 *This year, 120,800 people participated in the high school entrance examination. This number is for the entire city of Chengdu, including 20 districts, high-tech zone and Tianfu new district. A month ago, the Education Bureau announced the 2019 high school enrollment plan. The number of enrollment will be increased, indicating that there is a greater chance of going to high school... The plan of the central area (including 13 districts) is 43,015.*

Verdict: Refuted; **Domain**: Society

Challenges: Evidence Collection; Numerical Reasoning

Table 1: An example from CHEF (Chinese is translated into English). The claim is refuted by the evidence, which are sentences retrieved (highlighted) from the document. For brevity, only the relevant snippet of the document is shown.

Chen et al., 2020). There only exists a handful of non-English datasets for verifying real-world claims. However, these datasets are either small in size (Baly et al., 2018), or designed for multilingual systems (Gupta and Srikumar, 2021). On the other hand, Khouja (2020) and Nørregaard and Derczynski (2021) created claims by paraphrasing sentences from non-English articles, but synthetic claims cannot replace real-world claims for training generally applicable fact-checking systems.

To bridge this gap, we introduce a CHinese dataset for Evidence-based Fact-checking (CHEF). CHEF includes claims that are not only relevant to the Chinese world, but also originally made in Chinese. It consists of 10,000 real-world claims, collected from 6 Chinese fact-checking websites covered multiple domains and paired with annotated evidence. To ensure annotation consistency,

---

Dataset	Natural	Domain	#Claims	Language	Evidence			
					Type	Source	Retrieved	Annotated
FEVER (Thorne et al., 2018)	✗	Multiple	185,445	English	Text	Wiki	✓	✓
HOVER (Jiang et al., 2020)	✗	Multiple	26,171	English	Text	Wiki	✓	✓
TabFact (Chen et al., 2020)	✗	Multiple	92,283	English	Table	Wiki	✗	✓
InfoTabs (Gupta et al., 2020)	✗	Multiple	23,738	English	Table	Wiki	✗	✓
ANT (Khouja, 2020)	✗	Multiple	4,547	Arabic	✗	✗	✗	✗
VitaminC (Schuster et al., 2021)	✗	Multiple	488,904	English	Text	Wiki	✗	✓
DanFEVER (Nørregaard and Derczynski, 2021)	✗	Multiple	6,407	Danish	Text	Wiki	✓	✓
FEVEROUS (Aly et al., 2021)	✗	Multiple	87,026	English	Text/Table	Wiki	✓	✓
PolitiFact (Vlachos and Riedel, 2014)	✓	Politics	106	English	Meta/Text	FC	✗	✗
PunditFact (Rashkin et al., 2017)	✓	Multiple	4,361	English	✗	✗	✗	✗
Liar (Wang, 2017)	✓	Multiple	12,836	English	Meta	FC	✗	✗
Verify (Baly et al., 2018)	✓	Politics	422	Mul(2)	Text	Internet	✓	✗
MultiFC (Augenstein et al., 2019)	✓	Multiple	36,534	English	Meta/Text	Internet	✓	✗
Snopes (Hanselowski et al., 2019)	✓	Multiple	6,422	English	Text	FC	✗	✗
SciFact (Wadden et al., 2020)	✓	Science	1,409	English	Text	Paper	✗	✗
PUBHEALTH (Kotonya and Toni, 2020b)	✓	Health	11,832	English	Text	FC	✗	✗
AnswerFact (Zhang et al., 2020)	✓	Product	60,864	English	Meta/Text	Amazon	✓	✗
FakeCovid (Shahi and Nandini, 2020)	✓	Health	5,182	Mul(3)	✗	✗	✗	✗
XFact (Gupta and Srikumar, 2021)	✓	Multiple	31,189	Mul(25)	Meta/Text	Internet	✓	✗
CHEF	✓	Multiple	10,000	Chinese	Meta/Text	Internet	✓	✓

Table 2: Comparisons of fact-checking datasets. Type in the header means the type of evidence used, which can be text, metadata or both. Source means where the evidence are collected from, such as Wikipedia (Wiki), fact-checking websites (FC). Retrieved denotes if the evidence is given or retrieved from the source. Annotated means whether the evidence is manually annotated. Verify, FakeCovid, XFact contain claims in multiple languages, but Chinese claims are not included.

we developed suitable guidelines and performed data validation[2]. We shared some of the insights obtained during the annotation process that we hope will be beneficial to other non-English annotation efforts. Table 1 shows an instance from CHEF. In order to verify the claim, one needs to first retrieve the evidence sentences from related documents (e.g. government reports), then predict the veracity based on the evidence. After comparing the statistics of the entire city and central area, we can reach the verdict that the claim is refuted by evidence. To characterize the challenge of the dataset presented, we perform a thorough analysis and demonstrate the utility of the dataset by developing two types of baselines, including pipeline and joint systems. Our key contributions are summarized as follows:

1. We provide the first sizable multi-domain Chinese dataset for automated fact-checking. It consists of 10K real-world claims with manually annotated evidence sentences.

2. We further propose an approach that is able to model the evidence selection as a latent variable, which can be jointly trained with the veracity prediction module.

3. We develop several established baselines and

[2]The annotation guideline is provided in the appendix.

conduct a detailed analysis of the systems evaluated on the dataset, identifying challenges that need to be addressed in future research.

2 Background: Dataset Comparisons

In this section, we reviewed the existing fact-checking dataset as summarized in Table 2. Following Guo et al. (2022), we grouped the datasets into two categories: natural and synthetic. Natural datasets consist of real-world claims, while synthetic datasets contain claims created artificially by mutating sentences from Wikipedia articles.

2.1 Non-English Dataset

Existing efforts in the construction of non-English datasets are limited, both in scope and in size. Verify (Baly et al., 2018) contains 422 claims in Arabic and FakeCovid (Shahi and Nandini, 2020) has 3,306 non-English claims about COVID-19. Though XFact (Gupta et al., 2020) includes 31,189 claims in 25 languages, it mainly focuses on the multilingual setting, where the average number of instances per language is 1,248. More importantly, these datasets do not include annotated evidence. For example, XFact used search summaries returned by Google as evidence. On the other hand, Khouja (2020) and Nørregaard and Derczynski (2021) constructed synthetic datasets by mutating sentences from Arabic news and Danish

Wikipedia articles, respectively. Unlike previous efforts, CHEF consists of 10K real-world claims paired with annotated evidence.

There exist Chinese datasets focused on rumor detection (Ma et al., 2016; Zhang et al., 2021), which is classified into claim detection (Kotonya and Toni, 2020a; Guo et al., 2022), as it is based on language subjectivity and growth of readership (Qazvinian et al., 2011). A claim can be factual regardless of whether it is a rumour (Zubiaga et al., 2018). Unlike existing rumor detection datasets, CHEF focuses on factuality of the claim.

2.2 Evidence-Based Fact-Checking

Early efforts predicted the veracity solely based on the claims or with metadata (Rashkin et al., 2017; Wang, 2017), but relying on surface patterns of claims without considering the state of the world fails to identify well-presented misinformation (Schuster et al., 2020). Therefore, synthetic datasets (Thorne et al., 2018; Jiang et al., 2020; Aly et al., 2021) considered Wikipedia as the source of evidence and annotated the sentences supporting or refuting each claim. However, these efforts restricted world knowledge to a single source (i.e. Wikipedia), ignoring the challenge of retrieving evidence from heterogeneous sources on the Internet.

To address this, recent natural datasets (Augenstein et al., 2019; Gupta and Srikumar, 2021) used the summary snippets returned by Google as evidence. One key limitation of this approach is that summary snippets do not provide sufficient information to verify the claim. Gupta and Srikumar (2021) showed that only 45% of snippets provide sufficient information, while 83% of the full text from web pages provides sufficient evidence to determine veracity of the claim. To construct a better evidence-based dataset, we retrieve documents from web pages and manually select relevant evidence sentences from documents as evidence. Such a design makes CHEF suitable to train fact-checking systems that can extract evidence from web-sources and validate real-world claims based on evidence found on the Internet.

3 Dataset Construction

CHEF is constructed in four stages: data collection, claim labeling, evidence retrieval and data validation. Data collection selects sources, crawls claims and associated metadata. Claim labeling identifies claims from fact-checking articles and

Website	Domain	URL	Total
Piyao	Multiple	www.piyao.org.cn	3,741
TFC	Multiple	tfc-taiwan.org.tw	1,759
Mygopen	Multiple	www.mygopen.com	1,654
Jiaozhen	Multiple	vp.fact.qq.com	157
Cnews	Multiple	m.chinanews.com	2,689
Total	Multiple	-	10,000

Table 3: Statistics of data sources. Piyao, TFC, My-GoPen and Jiaozhen are fact-checking websites. Cnews is a news website.

assigns the veracity labels of claims based on the article. Evidence retrieval collects documents from the Internet and selects the most relevant sentences as evidence. Data validation controls the annotation quality. The annotation team has 25 members, 5 of them are only involved in data validation. All annotators are native Chinese speakers. To ensure the annotation quality, they were trained by the authors and went through several pilot annotations.

3.1 Data Collection

We crawled all active Chinese fact-checking websites listed by Duke Reporters[3]. However, most claims fact-checked by the fact-checkers are non-factual, solely relying on such claims will lead to an imbalance dataset. Therefore, we followed Kotonya and Toni (2020b) by crawling articles from the news review site. As shown in Table 3, this resulted in 5 websites in total. From each website, we crawled the full text of the article and corresponding metadata (e.g. author, domain, URL publication date). Totally, we crawled 14,770 fact-checking and news articles. There exists a number of crawling issues, such as the article could not be retrieved, or the content is not textual. We removed such instances. Next, we checked the dataset for duplications. Upon manual inspections, this was mainly due to them appearing on different websites. All duplications would be in the training split of the dataset, so that the model would not have an unfair advantage. As shown in Figure 5, claims cover multiple domains, including politics, public health, science, society and culture. More than 36% of claims belong to public health domain, as many fact-checking articles focused on countering misinformation related to COVID-19. The society domain has second most claims, which involves social events that are closely related to people's daily lives.

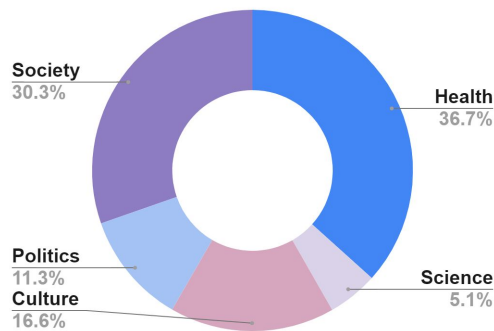---

[3] www.reporterslab.org/fact-checking/

Figure 1: Distributions of domains. Each instance is categorized into five different domains.

3.2 Claim Labeling

The major challenge of constructing a non-English dataset is extracting a claim and its veracity from a fact-checking article usually requires human efforts. Unlike fact-checking articles in English, many non-English articles (e.g. Chinese, Hindi, Filipino) do not explicitly give the claim and assign the veracity. Therefore, extracting the claim, which can appear in the title, or anywhere in the article, requires manual efforts. Before labeling the claim, we need to extract them from the fact-checking articles. When performing claim extraction, annotators need to read the fact-checking article first, then identify the claim. They are encouraged to select sentences directly from the article. However, resulted sentences may not be complete, which means they do not provide enough context for fact-checking. One common case is that the sentence describing a fact often lacks the time stamp, or the location. For example, the claim "Price of pork increases dramatically due to the African swine fever." is factual in 2020 but non-factual in 2021. Therefore, annotators are asked to complete the claim by adding missing content to ensure the claim to be standalone for later verification[4]. Another issue is that Chinese fact-checkers tend to use rhetorical questions to express non-factual claims. To alleviate the bias that the factuality of a claim can be decided by its surface form, annotators are required to paraphrase the questions into declarative statements.

Next, annotators are required to label the extracted claims. English fact-checking articles often provide different truth-rating scales, such as false, mostly false and mixture, while many non-English counterparts do not have such taxonomies. Therefore, annotators need to label the extracted claim based on the understanding of the fact-checking
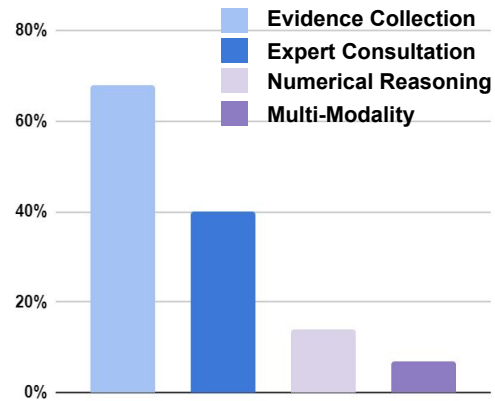

Figure 2: Distributions of challenges. Each instance can have multiple challenges. Evidence collection means finding relevant textual information from web-sources. Expert consultation collects information directly from relevant people. Numerical reasoning requires inference over numbers and multi-modality requires collect and infer over multi-modal evidence.

article. Journalism researchers showed that fine-grained labels are often assigned inconsistently due to subjectivity (Uscinski and Butler, 2013; Lim, 2018). Therefore, we chose to follow previous efforts (Thorne et al., 2018; Hanselowski et al., 2019) by adopting three types of labels: supported (SUP), refuted (REF) and not enough information (NEI), given the evidence. The distribution of labels in CHEF is shown in Table 4. CHEF consists of a majority of refuted claims, as the majority of fact-checking articles aim to debunk non-factual claims.

3.3 Evidence Retrieval

When verifying a claim, journalists first find information relating to the fact and evaluate it given the collected evidence. As shown in Figure 1, the biggest challenge of verifying a claim is to collect relevant evidence. In order to validate real-world claims, we chose to manually extract evidence from web-sources. We have two measures to ensure the reliability of the evidence. Firstly, we maintained a list of misinformation and disinformation websites, all search results from these websites will be filtered out. Secondly, we required the annotators to manually select evidence sentences from the search results. In order to collect evidence from the web-sources, we first submitted each claim as a query to the Google Search API by following Augenstein et al. (2019) and Gupta and Srikumar (2021). The ten most highly ranked search results are retrieved. For each result, we saved the search rank, URL, time stamp and document. Then we filtered out results from fact-checking websites to prevent the

---

[4]More annotation details are provided in appendix.

Split	SUP	REF	NEI	Total
Train	2,877	4,399	776	8,002
Dev	333	333	333	999
Test	333	333	333	999
Avg #Words in the Claim				28
Avg #Words in the Google Snippets				68
Avg #Words in the Evidence Sentences				126
Avg #Words in the Source Documents				3,691

Table 4: Dataset split sizes and statistics for CHEF.

answer from being trivially found. Next, annotators were asked to select sentences from the resulted documents. To maintain a balance between keeping relevant and removing irrelevant information, we followed Thorne et al. (2018) and Hanselowski et al. (2019) that up to five sentences were selected as evidence. Before deciding which sentences should be selected, annotators were required to answer auxiliary questions, such as "Whether selected sentences provide sufficient information for factual verification?" They were encouraged to select the five most relevant sentences, but they were allowed to pick less when relevant sentences are not available. A small fraction (5.6%) of instances do not have any relevant evidence, and we chose to discard them.

3.4 Data Validation

To ensure the annotation consistency, we conducted an additional 5-way inter-annotator agreement and manual validation. For inter-annotator agreement, we randomly selected 3% (n = 310) of claims to be annotated by 5 annotators. We calculated the Fleiss K score (Fleiss, 1971) to be 0.74, which is comparable with 0.68 reported in Thorne et al. (2018) and 0.70 in Hanselowski et al. (2019). In order to verify if evidence sentences provide sufficient information, we chose another 310 instances. The second group of annotators were required to assign the labels based on the evidence sentences. We found that 88.7% of the instances were labeled correctly and 83.6% of them provided sufficient information to determine the veracity. Finally, as shown in Table 4, we partitioned the dataset CHEF into training, development and test sets. Our development and test sets have balanced class distributions. Each claim is paired with Google snippets, evidence sentences and source documents.

4 Baseline Systems

Unlike previous natural datasets, CHEF requires the system to first retrieve the evidence sentences from the documents, then predict the veracity based on the evidence. Therefore, we design two types of baselines: pipeline and joint systems.

4.1 Pipeline System

The pipeline system treats evidence retrieval and veracity prediction as two independent steps.

4.1.1 Evidence Retrieval

Given the claim and documents, this step aims to select the most relevant sentences from documents as evidence, which can be viewed as a ranking problem. Thus, we adopt the following models:

Surface Ranker Following retrieval models designed for synthetic datasets (Thorne et al., 2018; Jiang et al., 2020; Aly et al., 2021), We use TF-IDF to sort the most similar sentences first and tune a cut-off using validation accuracy on the dev set.

Semantic Ranker Inspired by Nie et al. (2019) and Liu et al. (2020), we choose semantic matching based on BERT (Devlin et al., 2019) pre-trained on Chinese corpus (Wolf et al., 2020). The cosine similarity scores between the embedding of the claim and the embeddings of other sentences in the document are used for ranking.

Hybrid Ranker Since semantic encoding is complementary to surface form matching, they can be combined for better ranking. Following Shaar et al. (2020), we use the rankSVM, based on the feature sets of rankings returned by TF-IDF and the similarity scores computed with BERT.

Google Snippets As discussed in Section 2, existing natural datasets (Augenstein et al., 2019; Gupta and Srikumar, 2021) do not require the system to retrieve the evidence sentences from the documents. Instead, they used summary snippets returned by the Google Search Engine as evidence. We also include this type of evidence for comparisons.

4.1.2 Veracity Prediction

After retrieving the evidence sentences, veracity prediction aims to predict the label of the given claim. We implement the following classifiers:

BERT-Based Model Following Jiang et al. (2020) and Schuster et al. (2021), we use a multi-layer perceptron with embeddings from BERT as

the classifier. The embeddings of claim and retrieved evidence are concatenated as the input. The model performs the classification based on the output representation of the `CLS` token.

Attention-Based Model Following Gupta and Srikumar (2021), we first extract the output embedding of the `CLS` token of each selected evidence and calculate relevance weights with the claim through dot product attention. Then we feed the concatenated claim and weighted evidence into the BERT-based classifier.

Graph-Based Model Recent efforts (Zhou et al., 2019; Zhong et al., 2020; Liu et al., 2020) showed that graphs help to capture richer interactions among multiple evidence for fact-checking. We adopt the Kernel Graph Attention Network (Liu et al., 2020) for veracity prediction. The evidence graph is constructed based on the claim and evidence sentences, then node and edge kernels are used to conduct fine-grained evidence propagation. The updated node representations are used to calculate the claim label probability.

4.2 Joint System

Evidence retrieval in the pipeline system could not solicit optimization feedback obtained from veracity prediction. In order to optimize two steps jointly, we proposed to model the evidence retrieval as a latent variable. The joint system contains two modules: a latent retriever and a classifier. For the classifier, we used the same models described in Section 4.1. Latent retriever labels each sentence in the documents with a binary mask. Sentences labeled with 1 are selected as the evidence, while sentences labeled with 0 will be neglected.

4.2.1 Latent Retriever

We built the latent retriever based on the Hard Kumaraswamy distribution (Bastings et al., 2019), which gives support to binary outcomes and allows for reparameterized gradient estimates[5]. We first stretch the Kumaraswamy distribution (Kumaraswamy, 1980) to include 0 and 1 by the support of open interval (l, r) where $l < 0$ and $r > 1$, defined as $K^{'} \sim \mathrm{Kuma}(a, b, l, r)$ with CDF:

$$F_{K^{'}}(k^{'}; a, b, l, r) = F_{K^{'}}((k^{'} - l)/(r - l); a, b) \tag{1}$$

A sigmoid function $k^{''} = \min(1, \max(0, k^{'}))$ is used to rectify random variables into the closed interval $[0, 1]$, denoted by $K^{''} \sim \mathrm{HardKuma}(a, b, l, r)$ and $k^{''} = s(u; a, b, l, r)$ for short. Note that we map all negative values $k^{'} \in (l, 0]$ into $k^{''} = 0$ and $k^{'} \in [1, r)$ into $k^{''} = 1$ deterministically, so the sets whose masses under $\mathrm{Kuma}(k^{'} | a, b, l, r)$ are available in the closed form:

$$\begin{aligned} \mathbb{P}(K^{''} = 0) &= F_K(\frac{-l}{r - l}; a, b) \\ \mathbb{P}(K^{''} = 1) &= 1 - F_K(\frac{1 - l}{r - l}; a, b) \end{aligned} \tag{2}$$

Given source documents D, the latent retriever selects relevant sentences as evidence that can be used to predict the veracity for the claim c. For the i-th sentence $x_i \in D$, we obtain the sentence-level embedding $\mathbf{h}_i$ based on a BERT encoder by using the `CLS` token. Then we can calculate the latent selector $k_i^{''}$ by:

$$\begin{aligned} k_i^{''} &= s(u_i; a_i, b_i, l, r), \\ a_i = f_a(\mathbf{h}_i; \phi_a) \quad b_i &= f_b(\mathbf{h}_i; \phi_b) \quad u_i \sim \mathcal{U}(0, 1) \end{aligned} \tag{3}$$

where $f_a(\cdot; \phi_a)$ and $f_b(\cdot; \phi_b)$ are feed-forward networks with softplus outputs a_i and b_i. $s(\cdot)$ turns the uniform sample u_i into the latent selector $k_i^{''}$. Next, we use the sampled $k_i^{''}$ to modulate inputs to the classifier for veracity prediction:

$$f_F\left(k_i^{''} \cdot \mathbf{h}_i, \mathbf{c}; \theta_F\right) \tag{4}$$

where $\mathbf{c} = f_{\theta^{'}}(c)$ denotes the embedding for the given claim c obtained by using the `CLS` token through a BERT encoder. $f_F(\cdot; \theta_F)$ represents the classifier (e.g. graph-based model). The joint system can be optimized by gradient estimates of $\mathcal{E}(\phi, \theta)$ via Monte Carlo sampling from:

$$\mathcal{E}(\phi, \theta) = \mathbb{E}_{\mathcal{U}(0,1)}\left[\log P\left(y \mid X, s_\phi(u, X), \theta\right)\right] \tag{5}$$

where y is the label of veracity and $k_i^{''} = s_\phi(u, X)$ abbreviate the transformation from uniform samples to $\mathrm{HardKuma}$ samples.

4.2.2 More Baselines

Apart from the proposed system, we include following baselines for comprehensive comparisons:

Reinforce Instead of using gradient-based training, we follow Lei et al. (2016) by assigning a binary Bernoulli variable to each evidence sentence.

Because gradients do not flow through discrete samples, the evidence retriever is optimized using RE-INFORCE (Williams, 1992). A L_0 regularizer is used to impose sparsity-inducing penalties.

Multi-task We also adopt the multi-task learning method proposed by Yin and Roth (2018), which is the state-of-the-art joint model on FEVER dataset (Thorne et al., 2018). The model predicts a binary vector that indicates the subset of sentences as evidence, and a one-hot vector indicates the veracity of the claim. The overall training loss is the sum of these two prediction losses.

5 Experiments and Analyses

5.1 Experimental Setup

Following Augenstein et al. (2019), we computed the Micro F1 and Macro F1 as the evaluation metric. We further reported the mean F1 score and standard deviation by using 5 models from independent runs. For the pipeline system, we used 6 different evidence settings, including evidence sentences retrieved by surface ranker, semantic ranker and hybrid ranker, Google snippets, gold evidence and without using any evidence. For the joint system, we used 2 types of evidence: Google snippets and source documents, where the latent retriever can select sentences from. We used three classifiers for both systems, BERT-based (Schuster et al., 2021), attention-based (Gupta and Srikumar, 2021) and graph-based models (Liu et al., 2020).

The hyper-parameters are chosen based on the development set. In the evidence retrieval step of the pipeline system, we set the retrieved evidence obtained from TF-IDF to be more than 5 words for surface ranker. We use the BERT default tokenizer with max-length as 256 to preprocess data for semantic ranker. We use the default parameters in sklearn.svm.LinearSVC with RBF kernel for hybrid ranker.

In the veracity predication step of the pipeline system, we use the BERT default tokenizer with max-length as 256 and pretrained BERT-base-Chinese as the initial parameter to encode claim and evidence[6]. For BERT-based model, the fully connected network for classification is defined with layer dimensions of h_R-h_R/2-verification_labels, where $h_R = 768$. We use BertAdam (Devlin et al., 2019) with $5e{-}6$ learning rate, warmup with 0.1 to optimize the cross entropy loss and set the batch

size as 16. For attention-based model, we use BertAdam with $2e{-}5$ learning rate, warmup with 0.1 to optimize the cross entropy loss and set the batch size as 8. For graph-based model, we use BertAdam with $5e{-}5$ learning rate, warmup with 0.1, batch size with 16, dropout with 0.6 and kernel size with 21.

For the joint system, we use Adam (Kingma and Ba, 2015) with $5e{-}5$ learning rate, learning rate decay with 0.5 to optimize the cross entropy loss. We set the batch size as 32. The fully-connected networks $f_a(\cdot;\phi_a)$ and $f_b(\cdot;\phi_b)$ for two parameters a_i and b_i are defined with layer dimensions of $h_R = 768$. We set the dropout rate as 0.5.

5.2 Main Results

Pipeline System: According to Table 5, pipeline systems with evidence including Google snippets, sentences returned by rankers and gold evidence consistently outperform systems without using evidence. These results confirm that evidence plays an important role in verifying real-world claims. On the other hand, systems with retrieved sentences achieve higher scores than systems with Google snippets. Specifically, systems with gold evidence significantly outperform the ones with Google snippets, indicating information that is necessary for verification is missing in the snippets. Moreover, systems with retrieved evidence are more robust in terms of standard deviation. We hypothesize the reason is that irrelevant information is presented in the snippets. When comparing with different rankers, we observed that using contextualized representations to measure the similarity (Semantic Ranker) is generally better than exact string match (Surface Ranker). However, there still exists a large performance gap between the pipeline system with semantic ranker and the system with gold evidence. One potential solution is to develop better retrieval models based on the supervision signal of gold evidence provided by CHEF. Given the evidence sentences, graph-based models tend to have higher scores than BERT-based and attention-based models, which shows the effectiveness of leveraging graph structure to synthesize multiple evidence.

Joint System: Similar to the pipeline systems, joint systems that retrieve evidence sentences from documents achieve better F1 scores than directly use the summary snippets. In order to verify real-world claims, it is necessary to train fact-checking systems that learn how to effectively retrieve evi-

System / Evidence		BERT-Based Model[1]		Attention-Based Model[2]		Graph-Based Model[3]	
		Micro F1	Macro F1	Micro F1	Macro F1	Micro F1	Macro F1
Pipeline	No Evidence	54.46±2.89	52.49±2.44	54.87±1.95	53.47±2.82	—	—
	Snippets	62.07±2.55	60.61±2.96	62.42±2.31	60.24±2.56	62.78±1.70	61.06±2.59
	Surface Ranker	63.17±1.67	61.47±2.02	63.77±1.89	62.65±2.32	64.58±1.45	61.46±1.72
	Semantic Ranker	**63.47±1.71**	**61.94±1.66**	**63.95±1.46**	**62.80±1.33**	**64.67±1.54**	62.28±1.50
	Hybrid Ranker	63.29±1.65	61.80±2.31	63.48±1.22	62.74±1.30	64.37±1.66	**62.58±1.43**
Joint	Reinforce[4] — Snippets	63.76±1.52	61.74±1.88	64.06±1.76	61.97±1.04	65.77±1.23	62.34±1.11
	Reinforce[4] — Documents	64.37±1.65	62.46±1.72	64.86±1.83	62.66±1.32	66.58±1.45	63.47±1.58
	Multi-task[5] — Snippets	62.78±1.41	61.98±2.59	64.43±1.72	61.58±1.34	66.21±1.57	63.15±1.46
	Multi-task[5] — Documents	65.02±1.46	63.12±1.78	65.45±1.59	62.94±2.03	67.46±1.72	64.31±1.81
	Latent — Snippets	64.45±1.68	62.52±2.23	65.73±1.75	63.44±1.68	67.81±1.74	64.34±1.57
	Latent — Documents	**66.77±1.43**	**64.65±1.74**	**67.62±1.48**	**64.81±1.26**	**69.12±1.13**	**65.26±1.67**
Pipeline	Gold Evidence	**78.99±0.82**	**77.62±1.02**	**79.18±1.07**	**78.36±1.40**	**79.84±1.24**	**78.47±1.17**

Schuster et al. (2021)[1], Gupta and Srikumar (2021)[2], Liu et al. (2020)[3], Lei et al. (2016)[4], Yin and Roth (2018)[5]

Table 5: Results of pipeline and joint systems on CHEF.

dence sentences from full documents on web pages. In addition, joint system outperforms pipeline system consistently with both Google snippets and source documents as inputs. For example, latent retriever with Google snippets are able to achieve an average 2.74% and 1.77% Micro/Macro F1 boost compared with the pipeline systems with the same type of evidence. We attribute the consistent improvement of joint system to the explicit feedback to the evidence retrieval via gradient estimation on veracity prediction. Another advantage of the joint system is that the latent evidence retriever is able to dynamically select relevant sentences from documents for each instance, while rankers return a fixed number of evidence.

Compared with the reinforce and multi-task methods, the proposed latent retriever achieves 1.41% and 1.98% higher F1 on average with Google snippets and source documents as inputs across various classifiers. When considering standard deviation, reinforce is less robust. We believe the main reason is that latent retriever facilitate training through differentiable binary variables, which leads to robust and generalized model that exhibits small variance over multiple runs.

5.3 Analysis and Discussion

In this section, we further provide fine-grained analyses for baseline systems on CHEF. For brevity, we abbreviate pipeline systems with Google Snippets, Surface Ranker, Semantic Ranker, Hybrid Ranker as **GS**, **Sur**, **Sem**, **Hyb**, while joint systems with Google snippets and source documents as inputs as **JG** and **JS**, respectively. All results are reported

#E	GS	Sur	Sem	Hyb	JG	JS
1	55.24	55.67	56.04	56.72	56.98	57.54
3	58.69	59.24	59.52	59.18	59.89	61.45
5	**60.61**	**61.47**	**61.94**	**61.80**	**62.12**	64.65
10	59.12	60.20	60.37	61.24	61.86	**64.73**
15	55.72	56.31	56.56	57.08	58.69	59.11

Table 6: Effects of evidence: Macro F1 scores on the test set are reported. #E indicates the number of evidence.

based on the BERT-based model. We further provide case study and error analysis on CHEF. Due to limited spaces, we attach them in the appendix.

Effect of Evidence: In Table 6, we varied the numbers of evidence retrieved and reported the Macro F1 on the test set. The fluctuation results indicate that both quantity and quality of retrieved evidence affect the performance. Using fewer evidence will lead to incomplete coverage, which may not provide sufficient information to verify the claims. On the other hand, incorporating more evidence may introduce irrelevant sentences thus propagate errors to veracity prediction. In general, systems with 5 evidence sentences achieves the best performance except the joint system with source documents as inputs. We believe the reason is that the latent retriever maintains a better balance between keeping relevant and removing irrelevant sentences, which helps to achieve higher scores with more evidence sentences.

Performance against Claim Length: We partitioned the test set into 4 classes (<10, 10-19, 20-29, ≥30) based on lengths of the claims and reported
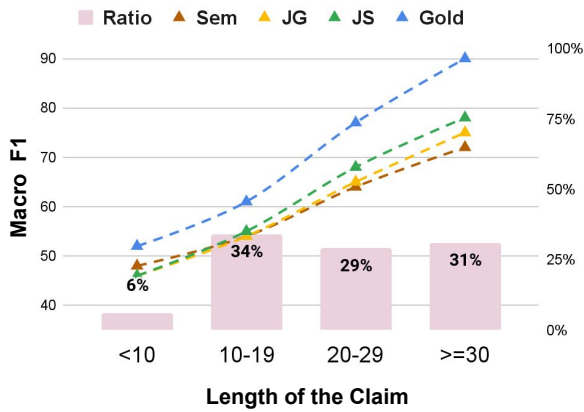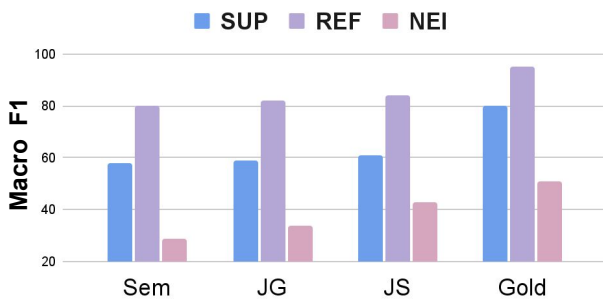
Figure 3: Comparisons against claim lengths.

Figure 4: Per-class results.

the Macro F1 score. For clarity, we choose the best reported pipeline system with semantic ranker to compare with joint systems. As shown in Figure 3, most claims are longer than 10 words. Performance of the systems on short claims (e.g. <10) is lower than other. One reason is that such claims do not contain sufficient information to retrieve evidence and to be verified, based on the observation that the performance of all the systems improve as the length of the claim increase. In general, the joint system outperforms the pipeline system against various claim lengths.

Performance against Classes and Domains: As CHEF is constructed based on real-world claims,
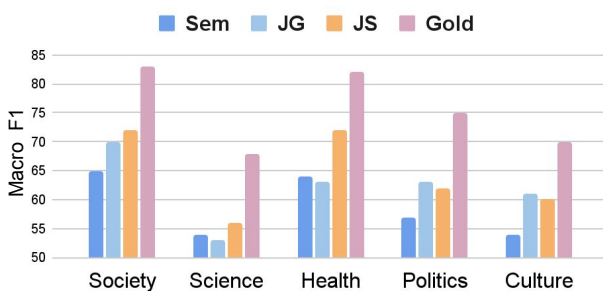

Figure 5: Per-domain results.

most of them are non-factual claims verified by fact-checking websites. Such an imbalance issue poses a challenge to the fact-checking system. Figure 4 shows the performance of models for different veracity labels. The scores of minor classes are much lower than the majority class. This reflects the difficulty of judging SUP and NEI. Informative evidence helps to alleviate this issue. For example, the pipeline system with gold evidence achieves significant improvement on predicting NEI labels when comparing with the system with semantic ranker. Figure 5 shows the performance of different domains. Claims from science, politics and culture domains have fewer training instances as most claims in the dataset focus on the society and public health topics. Again, retrieving informative evidence sentences (JS and Gold) from full documents is beneficial to this data sparsity issue.

6 Conclusion

We constructed the first Chinese dataset for evidence-based fact-checking. Further, we have discussed the annotation methods and shared some of the insights obtained that will be useful to other non-English annotation efforts. To evaluate the challenge CHEF presents, we have developed established baselines and conducted extensive experiments. We show that the task is challenging yet feasible. We believe that CHEF will provide a stimulating challenge for automatic fact-checking.

7 Acknowledgement

8 Ethical Consideration

Datasets have been collected in a manner that is consistent with the terms of use of any sources and the intellectual property. For each annotator, we compensate based on the number and quality of annotated sentences. More details of our datasets are depicted in Section 3.

3370

References

Bill Adair, Chengkai Li, Jun Yang, and Cong Yu. 2017. Progress toward "the holy grail": The continued quest to automate fact-checking. In *Proc. of the 2017 Computation+Journalism Symposium*.

Rami Aly, Zhijiang Guo, M. Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, O. Cocarascu, and Arpit Mittal. 2021. FEVEROUS: Fact Extraction and VERification over unstructured and structured information. In *NeurIPS*.

Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims. In *Proc. of EMNLP-IJCNLP*, pages 4685–4697, Hong Kong, China. Association for Computational Linguistics.

Ramy Baly, Mitra Mohtarami, James Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018. Integrating stance detection and fact checking in a unified corpus. In *Proc. of NAACL-HLT*, pages 21–27, New Orleans, Louisiana. Association for Computational Linguistics.

Jasmijn Bastings, Wilker Aziz, and Ivan Titov. 2019. Interpretable neural predictions with differentiable binary variables. In *Proc. of ACL*, pages 2963–2977. Association for Computational Linguistics.

Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang. 2020. Tabfact: A large-scale dataset for table-based fact verification. In *Proc. of ICLR*. OpenReview.net.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proc. of NAACL-HLT*, pages 4171–4186. Association for Computational Linguistics.

Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.

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 *Proc. of ACL-IJCNLP*, 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 *Proc. of ACL*, 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 *Proc. of CoNLL*, pages 493–503, Hong Kong, China. Association for 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 EMNLP*, pages 3441–3460, Online. Association for Computational Linguistics.

Jude Khouja. 2020. Stance prediction and claim verification: An Arabic perspective. In *Proc. of FEVER@ACL*, pages 8–17, Online. Association for Computational Linguistics.

Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proc. of ICLR*.

Neema Kotonya and Francesca Toni. 2020a. Explainable automated fact-checking: A survey. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 5430–5443. International Committee on Computational Linguistics.

Neema Kotonya and Francesca Toni. 2020b. Explainable automated fact-checking for public health claims. In *Proc. of EMNLP*, pages 7740–7754, Online. Association for Computational Linguistics.

Ponnambalam Kumaraswamy. 1980. A generalized probability density function for double-bounded random processes. *Journal of hydrology*, 46(1-2):79–88.

Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *Proc. of EMNLP*, pages 107–117.

Chloe Lim. 2018. Checking how fact-checkers check. *Research Politics*, 5.

Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2020. Fine-grained fact verification with kernel graph attention network. In *Proc. of ACL*, pages 7342–7351, Online. Association for Computational Linguistics.

Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J. Jansen, Kam-Fai Wong, and Meeyoung Cha. 2016. Detecting rumors from microblogs with recurrent neural networks. In *Proc. of IJCAI*, pages 3818–3824. IJCAI/AAAI Press.

Yixin Nie, Haonan Chen, and Mohit Bansal. 2019. Combining fact extraction and verification with neural semantic matching networks. In *Proc. of AAAI*.

Jeppe Nørregaard and Leon Derczynski. 2021. Danfever: claim verification dataset for danish. In *Proc. of the 23rd Nordic Conference on Computational Linguistics, NoDaLiDa 2021, Reykjavik, Iceland*

(Online), May 31 - June 2, 2021, pages 422–428. Linköping University Electronic Press, Sweden.

Vahed Qazvinian, Emily Rosengren, Dragomir R. Radev, and Q. Mei. 2011. Rumor has it: Identifying misinformation in microblogs. In *Proc. of EMNLP*.

Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proc. of EMNLP*, pages 2931–2937, Copenhagen, Denmark. Association for Computational Linguistics.

Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. Get your vitamin C! robust fact verification with contrastive evidence. In *Proc. of NAACL-HLT*, pages 624–643, Online. Association for Computational Linguistics.

Tal Schuster, Roei Schuster, Darsh J. Shah, and Regina Barzilay. 2020. The limitations of stylometry for detecting machine-generated fake news. *Computational Linguistics*, 46(2):499–510.

Shaden Shaar, Nikolay Babulkov, Giovanni Da San Martino, and Preslav Nakov. 2020. That is a known lie: Detecting previously fact-checked claims. In *Proc. of ACL*, pages 3607–3618. Association for Computational Linguistics.

Gautam Kishore Shahi and Durgesh Nandini. 2020. FakeCovid – a multilingual cross-domain fact check news dataset for covid-19. In *Proc. of WSM@AAAI*.

James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *Proc. of NAACL-HLT*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.

Joseph E Uscinski and Ryden W Butler. 2013. The epistemology of fact checking. *Critical Review*, 25(2):162–180.

Andreas Vlachos and S. Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proc. of LTCSS@ACL*.

David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. In *Proc. of EMNLP*, pages 7534–7550, Online. Association for Computational Linguistics.

William Yang Wang. 2017. "liar, liar pants on fire": A new benchmark dataset for fake news detection. In *Proc. of ACL*, pages 422–426, Vancouver, Canada. Association for Computational Linguistics.

Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proc. of EMNLP Demos*, pages 38–45. Association for Computational Linguistics.

Wenpeng Yin and Dan Roth. 2018. Twowingos: A two-wing optimization strategy for evidential claim verification. In *Proc. of EMNLP*, pages 105–114. Association for Computational Linguistics.

Wenxuan Zhang, Yang Deng, Jing Ma, and Wai Lam. 2020. AnswerFact: Fact checking in product question answering. In *Proc. of EMNLP*, pages 2407–2417, Online. Association for Computational Linguistics.

Xueyao Zhang, Juan Cao, Xirong Li, Qiang Sheng, Lei Zhong, and Kai Shu. 2021. Mining dual emotion for fake news detection. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, pages 3465–3476. ACM / IW3C2.

Wanjun Zhong, Jingjing Xu, Duyu Tang, Zenan Xu, Nan Duan, Ming Zhou, Jiahai Wang, and Jian Yin. 2020. Reasoning over semantic-level graph for fact checking. In *Proc. of ACL*. Association for Computational Linguistics.

Jie Zhou, Xu Han, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2019. GEAR: Graph-based evidence aggregating and reasoning for fact verification. In *Proc. of ACL*, pages 892–901, Florence, Italy. Association for Computational Linguistics.

Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. Detection and resolution of rumours in social media: A survey. *ACM Comput. Surv.*, 51(2):32:1–32:36.

A Supplement Materials

A.1 CHEF annotation guidelines

A.1.1 声明抽取和规范化的指引 Guidelines for claim extraction and normalization:

标注者首先需要认真阅读事实验证的文章，然后使用一到两句话来概括这篇文章描述的事件作为声明。请标注者直接使用文章中的句子作为声明，比如文章的标题，或者首段的前几句话都可能是这篇文章需要验证的事实。如果没法抽取出相应的句子，可以使用自己的语言总结文章来撰写声明。在撰写声明的时候，有以下注意事项：

- 每个声明必须完整。

- 每个声明不应该存在事实验证偏差。

- 每个声明不应该存在信息泄露。

请仔细阅读以下详细指引和相应的规范化例子：

- 声明中描述的事件缺乏必要的细节,比如：时间和地点。标注者需要加上这些细节让声明完整，才能被验证。比如：今年共有12.08万人参加中考，但招生计划只有4.3万。需要改写声明为：2019年，共有12.08万人参加成都中考，但招生计划只有4.3万。

- 声明中存在特殊符号，需要去除特设符号避免声明中存在偏差。比如："纯天然"喷雾一喷"秒睡"。需要去除句子中的""，因为这些特殊符号隐晦地表达了这个声明其实是不实的。模型可以通过特殊符号直接判断一个声明的真实性。这个句子需要改写为：纯天然喷雾一喷秒睡。，这可以避免由于声明中的特殊符号""带来事实验证偏差。

- 声明中存在信息泄露，需要去除直接指出声明真实性的相关词语。比如：谣言！纯天然喷雾一喷秒睡。句子中使用的"谣言！"已经直接指出该声明是不实的，造成了信息泄露。需要改写声明为：纯天然喷雾一喷秒睡。不能在声明中出现诸如："谣言"、"错误"，"骗局"等信息泄露词。

- 声明中的反问句需要被改写为陈述句，由于采用反问句形式的声明大部分都是不实的，反问句的形式会造成数据集偏差。比如：别人打了新冠疫苗，我们就可以不打新冠疫苗吗？需要被改写为：别人打了新冠疫苗，我们就可以不打新冠疫苗。

- 不陈述事实的声明需要被丢弃。有两大类的声明是无法进行事实验证的。第一大类为表示推测的声明，比如：明年深圳房价会上涨。第二大类为表示个人意见的声明，比如：我认为特朗普应该连任。

- 声明中如果包含多个声明，需要拆分为多个声明逐一验证。比如：关于新冠疫苗接种的两个事实：第一，别人打了新冠疫苗，自己就可以不打新冠疫苗。其次，新冠疫苗只需要打一针就能具备新冠病毒防护能力。这个声明包括了两个子声明，需要被拆分为：别人打了新冠疫苗，我就可以不打新冠疫苗。第二个声明为：新冠疫苗只需要打一针就能具备新冠病毒防护能力。

The annotator first needs to read the fact-checking article carefully, and then use one or two sentences to summarize the event described in this article as a claim. The annotator is encouraged to directly extract the sentences in the article as the claim, such as the title of the article, or the first few sentences in the first paragraph. If the annotator cannot find the sentence that can serve as a claim, you can use your own language to write the claim. When extracting the claim, there are the following considerations:

- Each claim must be complete.

- For each claim, explicit bias should be removed.

- Each claim should not have information leakage.

Please read the following detailed guidelines and corresponding normalized examples carefully:

- If the event described in the claim lacks important details, such as time and location, annotator need to add these necessary metadata to make the claim complete before it can be verified. For example, a total of 120,800 people took the entrance examination this year, but the enrollment plan is only 43,000. The claim needs to be rewritten as follows: In 2019, a total of 120,800 people participated in the Chengdu high school entrance examination, but the enrollment plan was only 43,000.

- If there exist special symbols in the claim, such symbols that may lead to bias for claim verification should be removed. For example: "Natural spray" helps you "sleep instantly".

Quotation marks should be removed in the sentence, as these special symbols implicitly indicates such a claim is non-factual. The model can predict the veracity simply based on the special symbols in the claim. This claim needs to be rewritten as: Natural spray helps you sleep instantly.

- Claims contains words that will lead to information leakage should be removed. For example: Rumors! Natural spray helps you sleep instantly. The word "Rumor!" in the claim directly pointed out that the claim is nonfactual, causing information leakage. The word "Rumor!" should be removed. Do not include information leaking words such as "rumors", "errors", "scams", etc. in the claim.

- Claims used rhetorical questions need to be rewritten into declarative sentences. Since most of the claims in the form of rhetorical question are nonfactual, the form of the rhetorical question exhibits a bias in the dataset. For example: if someone else gets the COVID-19 vaccine, can we not get the vaccine? It needs to be rewritten as: if someone else gets a COVID-19 vaccine, we do not need to get the vaccine..

- Claims that do not related to factuality should be discarded. There are two major types of claims that cannot be verified. The first category is speculative claims, such as: Shenzhen housing prices will rise next year. The second category is claims expressing personal opinions, such as: I think Donald Trump should be the president.

- A claim contains multiple statements should be split into multiple claims to be verified one by one. For example, a claim stated that: First, if someone else gets the COVID-19 vaccine, you do not need to get one. Also, the COVID-19 vaccine only needs one shot to protect against the virus. This claim includes two sub-claims. It needs to be split into two claims.

A.1.2 声明标注的指引| **Guidelines for claim labeling**

标注者在抽取出和规范化声明之后，需要根据事实验证的文章给出的结论，给每个声明打上标签。我们提供了以下三种标签，请选择其中的一种。注意的是，对于大部分为真，部分为真，大部分为为假，部分为假和半真半假的情况，我们统一归类为信息不足：

- 支持，有充分证据表明这个声明是被证据所支持的。

- 反对，有充分证据表明这个声明是被证据所反对的。

- 信息不足，没有足够的证据表明这个声明是被支持还是反对。

After extracting and normalizing the claim, annotators needs to label each claim based on the conclusions of the fact-checking article. We provide the following three labels, please choose one of them. Note that for conclusions such as mostly true, partially true, mostly false, partial false and mixture, we consider them as not enough information:

- Supported, there is sufficient evidence to show that this claim is supported by the evidence.

- Refuted, there is sufficient evidence to show that this claim is refuted by the evidence.

- Not enough information, there is not enough evidence to show whether this claim is supported or refuted.

A.1.3 证据标注的指引| **Guidelines for evidence labelling**

标注者需要阅读规范化过后的声明，事实验证的文章还有搜集到的源文档。标注者首先需要理解文章的验证思路，再从源文档当中直接选择能够作为证据的句子。针对每个声明，标注者最少选择1个，最多选择5个相关的句子作为证据。在选择句子作为证据的时候有以下注意事项：

- 请标注者选择完整的句子，以句号为结束标志。

- 选择句子作为证据的条件是，在仅仅基于当前选中的句子作为证据的前提下，能够验证给定的声明。也就是说，选中的句子必须要提供给足够的信息来帮助判断声明的事实性。

- 如果出现多于5个句子能够作为证据的情况，选择你认为最相关的5个句子；或者能够形成推理逻辑链的句子；或者和事实验证文章推理过程最相似的句子。

- 如果出现源文档互相矛盾的情况，优先选择支持事实验证文章结论的文档，从中选择相关的句子作为证据。

- 如果提供的源文档并没有提供足够的证据来验证声明，请报告这条声明。

- 如果提供的源文档只包含和事实验证文章结论矛盾的证据，请报告这条声明。

The annotator needs to read the normalized claim, the fact-checking article and the collected documents. Annotators first need to understand the verification process in the fact-checking article, and then directly select sentences from the sources documents. These selected sentence are used as evidence to verify the claim. For each claim, annotators should select at least 1 and at most 5 relevant sentences as evidence. There are the following considerations when choosing sentences as evidence:

- Please select a complete sentence which ends with a period.

- When selecting sentences as evidence, the annotator should consider given the selected sentences if the given claim can be verified. In other words, the selected sentences must provide sufficient information to predict the factuality of the given claim.

- If there are more than 5 sentences that can be used as evidence, choose the 5 sentences that you think are the most relevant; or the sentences that can form a reasoning chain for verification; or the sentence that is most similar to the reasoning process of the fact-checking article.

- If there are conflicting source documents, the documents that support the conclusion of the fact-checking article should be considered, and the most relevant sentences in these documents are selected as evidence.

- If the source documents do not provide sufficient evidence to verify the statement, please report the claim.

- If the source documents only contains evidence that contradicts the conclusion of the fact-checking article, please report this claim.

A.1.4 数据验证的指引 Guidelines for data validation

给定一个声明和搜集到的证据句子，标注者需要根据证据去判断这个声明的真实性。如果标注者认为提供的声明缺失重要信息，或者是不可读的，请报告该条声明。我们提供了以下三种标签，请选择其中的一种：

- 支持，有充分证据表明这个声明是被证据所支持的。

- 反对，有充分证据表明这个声明是被证据所反对的。

- 信息不足，没有足够的证据表明这个声明是被支持还是反对。

Given a claim and the evidence sentences, the annotator needs to label the factuality of the claim based on the evidence. If the annotator believes that the given claim lacks important information or is unreadable, please report the claim. We provide three kinds of labels, please choose one of them:

- Supported, there is sufficient evidence to show that this claim is supported by the evidence.

- Refuted, there is sufficient evidence to show that this claim is refuted by the evidence.

- Not enough information, there is not enough evidence to show whether this claim is supported or refuted.

A.1.5 判断声明领域的指引 Guidelines for determining claim domain

标注者需要阅读声明，根据给出的五个领域判断声明属于哪个领域：

- 政治：主要是关于国际与国内政治等方面的声明。

- 公卫：主要是关于公共卫生方面的声明，比如有关新冠病毒，人体健康，食品安全等方面。

- 科学：主要是关于自然科学和工程技术等方面的声明。

- 文化：主要是关于历史，人文，娱乐，体育等方面的声明。

- 社会：主要是除了上述四类，社会生活方面的声明。

The annotator needs to read the claim and determine which domain the claim belongs to based on the five domains given:

- Politics: Claims mainly focus on international and domestic politics.

- Health: Claims mainly focus on public health, including topic related to COVID-19, health care, food safety, etc.

- Science: Claims mainly focus on natural science and technology.

- Culture: Claims mainly focus on history, humanities, entertainment, sports, etc.

- Society: Claims not related on the above four categories, and related to daily social life.

A.1.6 验证声明的挑战（多选） **Claim verification challenges (Multiple choice)**

对声明进行事实验证往往会遇到许多挑战，挑战可以分为以下四类，标注者需要阅读事实验证的文章，判断验证声明时会遇到哪些挑战：

- 证据搜集：通过搜集证据，比如找相关的新闻，论文，法律法规等来验证声明。

- 专家咨询：通过咨询专家或者相关人士，比如外交部发言人陈述，部委回复，记者采访等来验证声明。

- 数值推理：通过数值的比较，趋势的分析来验证声明。

- 多模态：通过除了文本外的其他证据，比如图片，视频，音频来验证声明。

Factual verification of a claim often encounters many challenges. The challenges are summarized into the following four categories. The annotator needs to read the fact-checking article to determine which challenges will be encountered in verifying the claim:

- Evidence Collection: Verify the claim by collecting evidence, such as finding relevant news, papers, laws and regulations, etc.

- Expert Consultation: Verify the claim by consulting experts or related people, such as statements by the spokesperson of the Ministry of Foreign Affairs, replies from ministries and commissions, interviews with reporters, etc.

- Numerical Reasoning: Verify the claim by numerical comparison, trend analysis, etc.

- Multi-Modality: Verify the claim with other evidence besides articles, such as pictures, videos, and audio.