

# Zero-Shot Fact-Checking with Semantic Triples and Knowledge Graphs

Zhangdie Yuan and Andreas Vlachos

Department of Computer Science and Technology

University of Cambridge

zy317, av308@cam.ac.uk

## Abstract

Despite progress in automated fact-checking, most systems require a significant amount of labeled training data, which is expensive. In this paper, we propose a novel zero-shot method, which instead of operating directly on the claim and evidence sentences, decomposes them into semantic triples augmented using external knowledge graphs, and uses large language models trained for natural language inference. This allows it to generalize to adversarial datasets and domains that supervised models require specific training data for. Our empirical results show that our approach outperforms previous zero-shot approaches on FEVER, FEVER-Symmetric, FEVER 2.0, and Climate-FEVER, while being comparable or better than supervised models on the adversarial and the out-of-domain datasets.

## 1 Introduction

Fact-checking is the task of assessing the truthfulness of a claim, and is well-studied across multiple disciplines. Traditionally, journalists perform such a task manually, which is time-consuming. More recently, automated fact-checking systems have become of interest due to the explosion of (mis)information on social media (Adair et al., 2017; Hassan et al., 2017). In the NLP community, fact-checking is typically defined as a task consisting of three stages: claim detection, evidence retrieval, and claim verification (Guo et al., 2022). In particular, verdict prediction assumes the evidence is retrieved from sources such as Wikipedia or the web, and aims to predict the verdict of a claim given the retrieved evidence, often as a three-way classification task (Thorne et al., 2018a): SUPPORTS, REFUTES, and NEI (NOT ENOUGH INFO).

Recent work (DeHaven and Scott, 2023) has achieved strong results on canonical datasets like FEVER (Thorne et al., 2018a), mostly relying on supervised approaches. However, concerns

have been expressed on whether these models learn language’s and the task’s nuances or merely leverage embedded biases and dataset idiosyncrasies. This argument (Gururangan et al., 2018; Poliak et al., 2018) gains empirical weight when such high-performing models are tested against adversarial fact-checking datasets such as FEVER-Symmetric (Schuster et al., 2019) and FEVER 2.0 (Thorne et al., 2019). Their underperformance (Thorne et al., 2018b) in these adversarial benchmarks exposes a lack of model robustness.

The narrative of this vulnerability extends to out-of-domain contexts as well. A pertinent example is the Climate-FEVER dataset—a platform for verifying real-world climate claims (Diggelmann et al., 2020). Supervised models, despite their commendable performance on the original FEVER dataset, suffer performance degradation when evaluated on Climate-FEVER. Additionally, earlier zero-shot fact-checking approaches (Pan et al., 2021; Wright et al., 2022) hinge on synthetic data creation for training purposes. While this data emanates from factual evidence, it largely adheres to the domain boundaries of the originating dataset. Such inherent domain confinement curtails the model’s capacity for broader generalization.

In this work we propose a zero-shot method utilizing semantic triples and knowledge graphs in conjunction with pretrained Natural Language Inference (NLI) models, and does not require training data for parameter learning.

In particular, we propose to extract triples from the claim and the evidence texts to form knowledge graphs and fill potential gaps in the evidence using a universal schema model (Riedel et al., 2013) on Wikidata and Wikipedia. Crucially, our method refrains from utilizing any annotated or synthetic training data, sidestepping the pitfalls of biases and dataset artifacts that can inadvertently be encoded into models. Additionally, by decomposing the original claim into triples, our method can harness

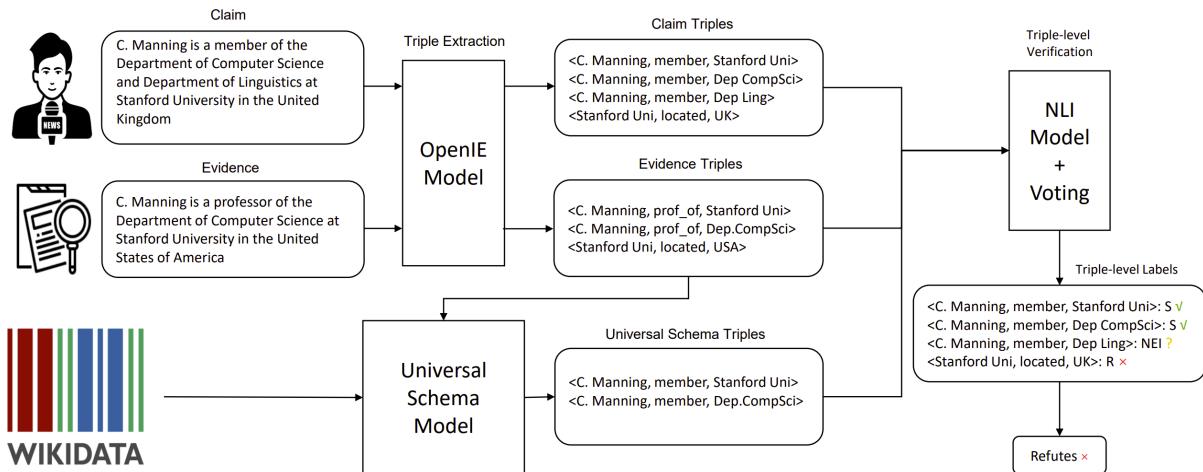


Figure 1: An overview of our zero-shot learning system. By harnessing Wikidata for training the universal schema model, incorporating on-demand training with evidence triples, and leveraging OpenIE for triple-level inference, our system achieves enhanced improvements. Label S stands for SUPPORTS and R stands for REFUTES.

the pre-trained NLI model’s strengths more effectively. Both design choices position our approach to exhibit greater robustness when subjected to adversarial and out-of-domain evaluations.

As shown in Figure 1, we follow a two-stage verification process: triple-level and claim-level. For triple-level, we employ NLI models pre-trained without further fine-tuning on any fact-checking training dataset, hence a zero-shot setting. For claim-level verification, we design a simple rule-based system relying on the triple verification. In Figure 1, the process involves extracting claim and evidence sentences to generate triples. Subsequently, the universal schema is applied to obtain additional triples. The NLI model is then employed to assign triple-level labels, resulting in 2 SUPPORTS, 1 REFUTES, and 1 NEI<sup>1</sup>. Finally, a rule-based system is utilized to derive claim-level verification. In this example, since one claim triple is refuted, the entire claim is considered refuted. Note that we are able to use the “gap” triples filled by the universal schema model to retrieve better evidence. For example, <Manning, member\_of, Stanford> is needed to verify the claim. However, such a triple is missing from the evidence triple extraction because the word *member* is not mentioned in the evidence. Instead, <Manning, professor\_of, Stanford> is extracted from evidence. Therefore, with the universal schema model, <Manning, member\_of, Stanford> will be assigned a high probability given <Manning, prof\_of, Stanford> is observed as evidence, and the gap is filled.

<sup>1</sup>In the context of our study, the NLI labels have been appropriately reconfigured to align with the FEVER labels.

We evaluate our approach on the FEVER (Thorne et al., 2018a), FEVER-Symmetric (Schuster et al., 2019), FEVER 2.0 (Thorne et al., 2019), and Climate-FEVER (Diggelmann et al., 2020) datasets. Our findings show that our system consistently outperforms zero-shot NLI model baselines by a margin of approximately 2.5 percentage points and beats the previous zero-shot approach by around 3 percentage points on FEVER-Symmetric. Notably, in contrast to state-of-the-art supervised methods (DeHaven and Scott, 2023), our approach exhibits robustness on both adversarial datasets. When evaluated on the out-of-domain Climate-FEVER dataset, our method outperforms the supervised method by a margin exceeding 10 percentage points.

## 2 Related Work

Recent advances in natural language processing have highlighted significant challenges associated with supervised learning models. A prominent concern is the models’ tendency to learn dataset-specific biases, often at the expense of genuine linguistic understanding. For instance, Schuster et al. (2019) demonstrated the effectiveness of a claim-only model that classifies each claim in isolation, without the need for associated evidence. The high performance achieved by their system over the baseline can be attributed to the idiosyncrasies inherent in the dataset’s construction. Similarly, Thorne et al. (2019) highlighted the vulnerability of several FEVER systems, observing significant performance declines under adversarial conditions

with simple rule-based perturbations. In other tasks such as NLI, previous works (Poliak et al., 2018; Gururangan et al., 2018) examined the susceptibility of neural models to such spurious correlations, revealing a troubling propensity for models to exploit unintended, data-specific heuristics. Taken together, these findings suggest that annotation artifacts within datasets contain discernible patterns. Such vulnerabilities underscore the necessity for more rigorous evaluation mechanisms, thus motivating the introduction of several adversarial fact-checking evaluation datasets (Guo et al., 2022).

Pan et al. (2021) presented the first work to investigate zero-shot fact verification, where they proposed a framework named Question Answering for Claim Generation (QACG). From any given evidence, QACG generates SUPPORTS, REFUTES, and NEI claims. A classifier is then trained using the generated claims instead of annotated claims, hence a zero-shot setting. To generate claims, QACG first produces QA pairs using a Question Generator fine-tuned on the processed SQuAD dataset (Zhou et al., 2018). Next, a QA-to-Claim Model is fine-tuned on the QA2D dataset (Demszky et al., 2018), which converts each QA pair into a declarative sentence. However, their experiments are limited, using only the gold evidence to evaluate various zero-shot methods, which is not practical in a real-world setting. Also, unlike their work, where training is still performed using the generated training data, our approach does not require any training for claim verification.

Knowledge graphs have long been investigated in NLP, where the first discussions of a graphical knowledge representation can date back to the 50s (Newell et al., 1959). Since then, many NLP researchers have tried to integrate knowledge graphs into various NLP tasks, notably language models with knowledge graphs (Nakashole and Mitchell, 2014; Logan IV et al., 2019; Liu et al., 2020; Wang et al., 2021a) and many downstream tasks such as question answering (Liu et al., 2020) and text classification (Hu et al., 2021). For fact-checking specifically, Ciampaglia et al. (2015) proposed to use knowledge graphs to verify simple natural language claims, considering fact-checking as a special case of link prediction. Their method uses the subject and object of the claim and then finds the shortest path between the two entities. If the claim is true, there should be such a shortest path (or an edge); otherwise, there should be no shortest path (nor edge). While the fact that a simple

notation	description
$\mathcal{C}$	claim
$\mathcal{E}$	evidence
$\mathcal{Y}$	claim-level label of $\mathcal{C}$
$\mathbb{C}$	set of triples extracted from $\mathcal{C}$
$\mathbb{E}$	set of triples extracted from $\mathcal{E}$
$c$	$c \in \mathbb{C}$
$e$	$e \in \mathbb{E}$
$y_e$	triple-level label of $c$ predicted by $e$
$y$	aggregated triple-level label of $c$

Table 1: Notations used in our fact-checking system

shortest-path computation can assess the truth of new claims is exciting, this work is limited because all the factual claims are automatically generated using triples. Therefore, it does not directly apply to recent human-generated fact-checking datasets such as FEVER, as claims in FEVER are much more complicated.

### 3 Methodology

As introduced, the verdict prediction step of claim verification is to predict a label  $\mathcal{Y} \in \{\text{SUPPORTS}, \text{REFUTES}, \text{NEI}\}$  given a claim  $\mathcal{C}$  and its corresponding evidence  $\mathcal{E}$ , indicating if  $\mathcal{C}$  is supported, refuted, or cannot be verified by  $\mathcal{E}$ . While we do not use any training data (manually or automatically labeled), we assume a human-annotated development set is available for fine-tuning hyperparameters of our system. In keeping with prior research, we use the same set of notations and extend it to include triples. Table 1 contains all the notations used in the methodology.

Figure 1 illustrates the structure of our system, which comprises three main steps: Triple Extraction, Triple-level Verification, and Claim-level Verification. Additionally, we have integrated an external component, the Universal Schema. This section provides comprehensive insights into each component, outlining their functionalities and operations.

#### 3.1 Triple Extraction

A semantic triple consists of three entities: the subject, the object, and the relation between them. We denote such a triple as  $\langle \text{subj}, \text{rel}, \text{obj} \rangle$  where all three entities are natural language words, phrases, or clauses, and no schema needs to be specified in advance. Extracting a set of triples from plain text is called open information extraction (Open IE) (Yates et al., 2007).

As illustrated in Figure 1, our system first employs an OpenIE tool to extract triples from claim  $\mathcal{C}$  and evidence  $\mathcal{E}$ , resulting in a set  $\mathbb{C}$  of claim triples and a set  $\mathbb{E}$  of evidence triples. Note that this step is back-traceable. For example, for any evidence triple, we can trace back which evidence sentence it comes from and which part of that sentence forms such evidence triple.

### 3.2 Triple-level Verification

Given a claim triple  $c$  and a set of evidence triples  $\mathbb{E}$ , triple-level verification predicts a label  $y \in \{\text{SUPPORTS}, \text{REFUTES}, \text{NEI}\}$  for this triple. Intuitively, any evidence triple may provide signals to predicting  $y$ . Therefore, given  $c$ , for each evidence triple  $e$  in  $\mathbb{E}$ , we utilise an NLI model to predict a label  $y_p$  with a softmax score as its probability. We linearise  $c$  and  $e$  and concatenate them with a special separation token:  $e$  [SEP]  $c$ . For example, let  $c = \langle \text{Barack Obama, was born in, USA} \rangle$  and  $e = \langle \text{Barack Obama, was born in, Hawaii} \rangle$ , the input of the NLI model therefore is, ‘Barack Obama was born in Hawaii [SEP] Barack Obama was born in USA’. Following previous work, we map NLP labels to fact-checking labels, namely ENTAILMENT to SUPPORTS, CONTRADICTION to REFUTES, and NEUTRAL to NEI.

To filter out less reliable triple-level labels  $y_p$ , we set up two thresholds for SUPPORTS and REFUTES as hyperparameters to cut off labels with low probabilities. The remaining labels are aggregated to reach a triple-level label  $y$  for the claim triple  $c$  using one of the following voting mechanisms:

**Max voting** takes the label with the overall highest probability as the triple-level label.

**Majority voting** ensures that the label with the most supporters (i.e., most frequent appearances) is the triple-level label.

**Weighted sampling** samples a label according to the highest probabilities of each label.

Note that if all labels are filtered out, the triple-level label is NEI because none of the evidence triple is reliable enough for this claim triple.

### 3.3 Claim-level Verification

For each claim triple  $c$  in  $\mathbb{C}$ , a triple-level label  $y$  is predicted by the previous step. The final step is to reach a claim-level label  $\mathcal{Y}$  from this set of triple-level labels using the following rule-based system also used by previous research (Stacey et al., 2022):

- If there exists a  $y$  that is REFUTES, then  $\mathcal{Y}$  is REFUTES.

- If no  $y$  is REFUTES and there exists a  $y$  that is NEI, then  $\mathcal{Y}$  is NEI.
- Otherwise,  $\mathcal{Y}$  is SUPPORTS.

### 3.4 Universal Schema

The challenge of integrating a knowledge graph with our system stems from the incompatibility between a pre-determined schema and the unrestricted textual information extracted from open sources. In response, we put forward a solution that involves the implementation of the universal schema (Riedel et al., 2013), which acts as an interface between pre-defined symbolic relations such as those found in knowledge graphs, and unconstrained textual relations such as those extracted by Open IE. Universal schema can be viewed as a matrix that represents the knowledge base, comprising pairs of entities and relations.

Notably, the original knowledge graph dataset employed in previous research on Universal Schema, namely Freebase (Bollacker et al., 2008), is no longer maintained. Therefore, we undertook the task of training a novel universal schema model, utilizing a more contemporary language model architecture and incorporating data from Wikidata (Vrandečić and Krötzsch, 2014) and some corresponding texts from Wikipedia.

**Task Definition.** Consistent with Riedel et al. (2013), a fact, or relation instance, is denoted by the pair  $rel$  and  $\langle subj, obj \rangle$ . The goal of a universal schema model is, by definition, to estimate, for a given relation  $rel$  and a given tuple  $\langle subj, obj \rangle$ , the probability  $p(y_{rel, \langle subj, obj \rangle} = 1)$  where the random variable  $y_{rel, \langle subj, obj \rangle}$  represents if  $\langle subj, obj \rangle$  is in relation  $rel$ . In the context of our fact-checking scenario, we leverage these  $\langle subj, rel, obj \rangle$  triples to complete missing information.

**Objective.** To train our model, we adopt Bayesian Personalized Ranking (BPR) (Rendle et al., 2009). In this approach, observed true facts are assigned higher scores compared to both true and false unobserved facts. This scoring scheme serves as our optimization objective. Let  $\sigma$  denote the sigmoid function,  $\theta_{f+}$  denote the dot product of the latent representations of a positive ( $\theta_{f-}$  for negative) fact pair  $rel$  and  $\langle subj, obj \rangle$ , then the objective function is  $\text{Obj}_{f+, f-} = -\log(\sigma(\theta_{f+} - \theta_{f-}))$ .

**Integration.** Upon successful training of the universal schema model, it becomes feasible to predict the probability of a tuple being associated with a given relation. The integration of this component

into our fact-checking system involves utilizing the universal schema model to assign scores to potential triple candidates for a given set of claim triples  $\mathbb{C}$  and supporting evidence triple set  $\mathbb{E}$ . All possible combinations of relations in the  $\mathbb{C}$  and tuples in the  $\mathbb{E}$  are considered as triple candidates. The universal schema model is then used to compute the probability of each triple candidate being true. Similar to Section 3.2, a threshold is set to remove less reliable triple candidates. The triple candidates with a probability above the threshold are only utilized for triple-level verification if the available evidence triples are insufficient, i.e. when the label for the triple-level label  $y$  is NEI. In a manner akin to in-context learning, we also modify the Universal Schema model during the inference stage, upon encountering newly observed facts derived from evidence triples  $\mathbb{E}$ .

## 4 Implementation

**Evidence Retrieval.** To perform document level retrieval, we adopt the approach proposed by Hanselowski et al. (2018) For sentence level retrieval, we aim to demonstrate the effectiveness of our verification system without relying on any fact-checking training data. Therefore, we utilize traditional information retrieval techniques such as tf-idf weighting. In addition, we incorporate a semantic score as a weight factor, which is computed using the cosine similarity of embeddings generated by a neural model called stsb-roberta-base (Reimers and Gurevych, 2019).<sup>2</sup>

**OpenIE Model.** We utilized an AllenNLP reimplementation of a BiLSTM sequence prediction model initially proposed by Stanovsky et al. (2018) as our Open Information Extraction (OpenIE) tool. The model can recognize verbs as relations and add their corresponding subjects and objects as arguments when given a sentence as input. For instances with more than two arguments, the model produces a triple for each combination of subjects or objects. If a relation only has one argument, known as a unary relation, a placeholder is added to ensure consistency across all generated triples.

**NLI Model.** In our experiments, we evaluate the effectiveness of our system using both base size and large size pre-trained NLI models. The aim is to demonstrate that our system consistently outper-

forms the NLI baselines. In particular, we leverage the RoBERTa base and large models, which have been pretrained on the MNLI dataset. Both models follow the standard NLI format of taking a premise and a hypothesis as input in the format of "[premise] SEP [hypothesis]", where SEP denotes the special separation token. We adhere to this format throughout our experiments.

**Universal Schema Model.** We leverage Sentence-BERT (Reimers and Gurevych, 2019) to obtain sentence embeddings that serve as latent representations for both relations and tuples. This approach allows us to capture the semantic meaning of the sentences, which is essential for accurately representing the relations and tuples in our model. The pre-trained model "all-MiniLM-L6-v2"<sup>3</sup> is utilized in our study, which is based on MiniLM (Wang et al., 2020). This model has been pre-trained with a contrastive objective using diverse datasets containing sentence pairs. The cosine similarity is computed for each possible sentence pair within a batch, and cross-entropy loss is employed to compare these similarities with the true pairs.

## 5 Experimental Setup

**Dataset.** In our evaluation, we employ four benchmark datasets, FEVER, FEVER-Symmetric, FEVER 2.0, and Climate-FEVER. FEVER dataset (Thorne et al., 2018a) comprises 185,445 claims that are created by modifying sentences from Wikipedia, which are subsequently verified on Wikipedia without knowing the original sentence they were derived from. On the other hand, the FEVER-Symmetric dataset is introduced by Schuster et al. (2019) to address the biases identified in the original FEVER dataset. This dataset is constructed with a regularization procedure to downweigh the giveaway phrases that cause potential biases. Similarly, the FEVER 2.0 dataset (Thorne et al., 2019) comprises adversarial examples intentionally created by participants of the FEVER 2.0 shared task. The task required teams to generate claims specifically designed to challenge FEVER-trained models. From this dataset, we extracted all SUPPORTS and REFUTES claims, along with their corresponding gold evidence sentences, for our evaluation. The Climate-FEVER dataset (Diggelmann et al., 2020) is for

<sup>2</sup><https://huggingface.co/sentence-transformers/stsb-roberta-base>

<sup>3</sup><https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

verification of real-world climate change-related claims, excluding disputed claims.

To train our universal schema model, we utilize the Wikidata5m dataset (Wang et al., 2021b), which is a knowledge graph dataset comprising one million entities. This dataset is particularly suitable for our purposes, as it includes an aligned corpus, which we leverage in conjunction with the OpenIE tool to extract open-domain triples. In our approach, all triples from the Wikidata5m dataset and the extracted triples the aligned corpus are treated as positive samples. To generate negative samples based on a given positive sample, we utilize a randomized approach where we preserve the relation and generate arbitrary tuples that exist within the dataset. This approach allows us to create negative samples that differ from the positive samples while still being relevant to the original relation.

**Hyperparameters.** In our experiments, claim verification does not require model-level hyperparameter tuning since no training is involved. However, as outlined before, we have a small set of three thresholds to be adjusted: a threshold for the SUPPORTS label at the triple level, a threshold for the REFUTES label at the triple level, and a threshold for filtering out Universal Schema triple candidates.<sup>4</sup> We adjusted their values on the FEVER dataset and did not perform any further adjustments on FEVER-Symmetric, FEVER 2.0, or ClimateFEVER. This deliberate choice was made to test the robustness of our system in handling different datasets without relying on dataset-specific optimization, an advantage of zero-shot approaches.

Figure 2 illustrates the impact of thresholds on our system, and that optimizing them is relatively straightforward as the best settings are clustered in the region. It is worth noting that the optimal threshold for REFUTES is considerably higher compared to SUPPORTS, indicating that our system is more stringent in assigning a triple-level REFUTES label than SUPPORTS. This difference is justified by the fact that, as explained in Section 3.3, a single refuted claim triple is sufficient to refute the entire claim, therefore it helps being cautious when assigning a REFUTES label to claim triples.

<sup>4</sup>The specifics concerning the hyperparameters of the Universal Schema model can be found in Appendix A. Note that the aforementioned thresholds were identified by conducting a search with a fixed Universal Schema model.




Figure 2: The influence of thresholds on accuracy for the SUPPORTS (x-axis) and REFUTES (y-axis) with fixed threshold for Universal Schema triple candidates.

## 6 Results and Discussion

Our main results are presented in Table 2, where we compare our system’s performance against the current state-of-the-art system on FEVER, BEVERERS (DeHaven and Scott, 2023), the FEVER-trained entailment-predictor (Diggelmann et al., 2020), and QACG (Pan et al., 2021)<sup>5</sup>. Also, we conduct additional ablation experiments to demonstrate the robustness of our system by varying the weighting factor and voting mechanism.

**Finding 1** *Our zero-shot approach exhibited enhanced robustness against adversarial perturbations and manifested notable out-of-domain effectiveness in contrast to supervised approaches.*

As shown in Table 2, our zero-shot method demonstrates greater resilience against adversarial attacks compared to supervised methods, providing a significant advantage in real-world scenarios where the presence of misinformation and deceptive tactics can impede the performance of fact-checking systems.

By abstaining from using training data, our approach intuitively circumvents these issues and offers a more robust approach to fact-checking. Specifically, despite exhibiting lower performance than supervised systems on the original FEVER dataset, our models achieved highly competitive scores on the FEVER-Symmetric dataset, trailing

<sup>5</sup>We made efforts to establish contact with the authors of QACG; however, our attempts to elicit a response were unsuccessful. Therefore, a direct comparison with their approach is not feasible, except for the FEVER-Symmetric dataset, where they reported performance under the same setting as ours.

Model	FEVER		FEVER-Symmetric	FEVER 2.0	Climate-FEVER	
	Accuracy	FEVER Score	Accuracy	Accuracy	Accuracy	F1
<i>Supervised</i>						
BEVERS	<b>80.24</b>	<b>77.70</b>	75.9	63.4	-	-
Diggelmann et al. (2020)	77.69	-	-	-	38.78	32.85
<i>Zero-shot</i>						
Random Guess	33.33	-	50.00	50.00	33.33	33.33
QACG	-	-	77.1	-	-	-
NLI Model (base)	36.07	31.65	51.74	49.68	-	-
NLI Model (large)	58.12	54.38	78.94	70.38	44.37	44.24
Our system (base)	38.56	33.73	56.62	53.28	-	-
Our system (large)	60.40	56.79	<b>79.78</b>	72.92	<b>46.71</b>	<b>45.71</b>
Our system (large) + USchema	<b>61.30</b>	<b>57.84</b>	<b>79.78</b>	<b>73.34</b>	<b>46.71</b>	<b>45.71</b>

Table 2: Main results on FEVER (S/R/NEI), FEVER-Symmetric (S/R), FEVER 2.0 (S/R) and Climate-FEVER (S/R/NEI). BEVERS is the current state-of-the-art system on FEVER and Diggelmann et al. (2020) is the entailment-predictor based on ALBERT (large-v2). The accuracy on FEVER-Symmetric, FEVER 2.0 and Climate-FEVER datasets was achieved without fine-tuning, demonstrating the models’ robustness.

Variant	$\Delta$ Accuracy
tf-idf	<b>+2.28</b>
Cosine similarity	+2.11
Max voting	+2.05
Majority voting	+2.13
Weighted sampling	+1.93

Table 3: Improvements of our system over baseline using different retrieval weighting factor and voting technique are steady.

the state-of-the-art by only approximately 2 percentage points. We attribute this positive outcome to our system’s utilization of NLI models, which already demonstrate outstanding performance on this adversarial dataset. The results obtained on the FEVER 2.0 dataset align with FEVER-Symmetric and further strengthen our conclusions.

On Climate-FEVER, the supervised approach delineated by Diggelmann et al. (2020) achieved an accuracy of 38.78% and an F1 score of 32.85%. In comparison, our introduced zero-shot methodology showcased enhanced results, achieving an accuracy rate of 46.71% and an F1 score of 45.71%, which demonstrated a notable generalization ability. These findings suggest that our zero-shot method offers a promising avenue for improved performance in out-of-domain tasks.

**Finding 2** *Our system, utilizing triple-level inference, consistently improves over the baseline results irrespective of the NLI model used.*

In our experiments, our approach was able to im-

Evidence	$\Delta$ Accuracy
Gold + Random	<b>+7.93</b>
Gold + Retrieved (tf-idf)	+3.74
Retrieved (tf-idf)	+2.28

Table 4: Improvements of our system over baseline using gold evidence vs. retrieved evidence.

prove the performance of both the base size and large size NLI models by approximately 2.5%. These consistent improvements suggest that our approach can continue to benefit from the ongoing progress: as more advanced models are being developed, our system is expected to demonstrate even greater accuracy and reliability.

In addition, we performed ablation experiments to investigate the impact of various weighting factors and different voting mechanisms, as outlined in Section 3. The results, presented in Table 3, demonstrate that our system’s improvements over the baseline NLI models in Table 2 are consistently observed across all variants, indicating the reliability and robustness of our approach.

Furthermore, we conducted experiments to evaluate the effect of evidence quality on claim verification, as presented in Table 4. The Gold + Random method involves using gold-standard evidence for SUPPORTS and REFUTES claims, while random evidence is used for NEI claims. The Gold + Retrieved method is similar, but uses retrieved evidence instead of random evidence for NEI claims while still utilizing gold-standard evidence for SUPPORTS and REFUTES claims. The results indicate

<b>Claim</b>	The Adventures of Pluto Nash was <b>reviewed</b> by Ron Underwood .
<b>Evidence Sentences</b>	1: The Adventures of Pluto Nash is a 2002 Australian-American science fiction action comedy film starring Eddie Murphy -LRB- in a dual role -RRB- and directed by Ron Underwood . 2: Ron “ Thunderwood ” Underwood is a musician and director from Phoenix , Arizona . ...
<b>Evidence Triples</b>	<a 2002 Australian - American science fiction action comedy film, starring, Eddie Murphy> ...
<b>USchema Triples</b>	<The Adventures of Pluto Nash, <b>directed</b> , by Ron Underwood> ...

Table 5: An example with Universal Schema triples. Due to space limitations, not all sentences and triples for this example are shown. The table focuses on the most critical ones that effectively demonstrate our points.

that the performance improvements of our system increases as the quality of evidence improves, suggesting that our zero-shot approach benefits more from less noisy evidence. This likely due to the fact that our system relies on a strict set of rules to classify claims, which may be more sensitive to the presence of noise in the evidence. Thus, our system is likely to benefit from the continued development of better evidence retrieval systems.

**Finding 3** *Employing the Universal Schema model provides marginal gains by bridging the gaps between extracted claim and evidence triples.*

The Universal Schema model, despite its modest gains, contributes to enhancing the overall performance of our fact-checking system. In manual analysis of the results we found that integrating the Universal Schema model helps our approach in handling claims involving mutual exclusivity, resulting in increased accuracy. Mutual exclusivity denotes a situation in which two or more events cannot coexist simultaneously. To illustrate this, let us consider the claim in Table 5 *The Adventures of Pluto Nash was reviewed by Ron Underwood*, initially classified as NOT ENOUGH INFORMATION (NEI) in the absence of the Universal Schema model. This misclassification originated from the complexity of the retrieved evidence, which presented a complex sentence implying that Ron Underwood directed the movie, thereby refuting the claim. However, extracting the relation needed as evidence *<The Adventures of Pluto Nash, directed by, Ron Underwood>* posed challenges so it was not extracted. Consequently, due to the absence of this critical triple, the model erroneously labeled the claim as NEI. By incorporating the Universal Schema model, our system successfully recovered the missing evidence triple, while also recognizing the inherent mutual exclusivity between assuming

both the director and reviewer roles for the same movie. As a result, using the Universal Schema model accurately predicted the REFUTES label.

We also observed that the Universal Schema model offers limited assistance when applied to the two adversarial datasets considered. This is due to the fact that, in both the FEVER-Symmetric and FEVER 2.0 setting, all the necessary evidence is provided, unlike real-world scenarios. Consequently, the value provided by the Universal Schema model, which primarily focuses on filling gaps, becomes minimal since no gaps exist in the presence of sufficient evidence.

## 7 Conclusion

We introduced a novel zero-shot fact-checking method, translating claims and evidence into semantic triples with external knowledge graphs. This method surpasses other zero-shot baselines, impressively without direct FEVER dataset training. Its resilience is evident, avoiding the typical performance dips seen in supervised models on adversarial datasets like FEVER-Symmetric and FEVER 2.0. Also on the Climate-FEVER dataset, our approach outshines even supervised counterparts, highlighting its generalization prowess. Augmented by pretrained NLI models, our system’s robustness is further emphasized. As future steps, we aim to hone model interpretability, examine diverse knowledge graphs, and test our method’s versatility on other fact-checking datasets.

## Limitations

While our novel zero-shot learning method for fact-checking with semantic triples and knowledge graphs has shown promising results, there are several limitations that must be noted.

Firstly, our method’s language capabilities have



been exclusively tested on the English language, which poses an inherent limitation. Though the method was not specifically designed and implemented for English, the experiments were solely conducted using English datasets. Consequently, the potential effectiveness of our approach with other languages remains unverified. Differences in linguistic features and semantic triple structures across languages might present unique challenges that we have yet to encounter or address.

Secondly, our approach relies heavily on Wikipedia as both the source for datasets used in evaluation and the basis for our knowledge graphs. While Wikipedia is a vast and continually updated source of knowledge, its use as the sole source of data introduces biases and limitations. Wikipedia’s content is predominantly generated by its user community, which can lead to the inclusion of inaccuracies, cultural biases, or omissions. This limitation might affect the fact-checking capabilities of our model, as the reliability of its responses are directly proportional to the quality and accuracy of the information within Wikipedia.

Additionally, the reliance on a single source for data and knowledge graphs constrains the method’s applicability in fact-checking scenarios where knowledge outside of Wikipedia’s domain is required. It may also lead to an overfitting issue, as the model might be overly tuned to Wikipedia’s style and structure, limiting its performance when applied to different or broader sources.

In future work, addressing these limitations by incorporating support for multiple languages and expanding the data sources beyond Wikipedia would be essential steps towards enhancing the effectiveness and generalizability of our approach.

## Ethics Statement

The use of fact-checking datasets and systems has become increasingly important in combatting misinformation, and as such, it is necessary to consider the ethical implications of their use. One of the key concerns in this regard is the potential for biases in these datasets. Such biases can arise from various sources, including the selection and interpretation of sources, the types of claims being fact-checked, and the demographic characteristics of the individuals involved. These biases have the potential to perpetuate stereotypes and reinforce existing power dynamics, and thus it is the responsibility of researchers to ensure that they use representative and

unbiased datasets to train and evaluate their models. Transparency regarding any potential biases in models is also essential, and steps must be taken to mitigate any negative impact. By addressing these ethical concerns, researchers can promote the integrity of fact-checking and contribute to a more informed and equitable public discourse.

## Acknowledgements

Zhangdie Yuan and Andreas Vlachos are both supported by the ERC grant AVeriTeC (GA 865958).

## References

- Bill Adair, Chengkai Li, Jun Yang, and Cong Yu. 2017. Progress toward “the holy grail”: The continued quest to automate fact-checking. In *Computation+ Journalism Symposium*, (September).
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: A collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’08, page 1247–1250, New York, NY, USA. Association for Computing Machinery.
- Giovanni Luca Ciampaglia, Prashant Shiralkar, Luis M Rocha, Johan Bollen, Filippo Menczer, and Alessandro Flammini. 2015. Computational fact checking from knowledge networks. *PloS one*, 10(6):e0128193.
- Mitchell DeHaven and Stephen Scott. 2023. [Bevers: A general, simple, and performant framework for automatic fact verification](#). *arXiv preprint arXiv:2303.16974*.
- Dorottya Demszky, Kelvin Guu, and Percy Liang. 2018. Transforming question answering datasets into natural language inference datasets. *arXiv preprint arXiv:1809.02922*.
- Thomas Diggelmann, Jordan Boyd-Graber, Jannis Buihan, Massimiliano Ciaramita, and Markus Leipold. 2020. [Climate-fever: A dataset for verification of real-world climate claims](#). *arXiv preprint arXiv:2012.00614*.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112,

- New Orleans, Louisiana. Association for Computational Linguistics.
- Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz, and Iryna Gurevych. 2018. **UKP-athene: Multi-sentence textual entailment for claim verification**. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 103–108, Brussels, Belgium. Association for Computational Linguistics.
- Naemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by claim-buster. In *Proceedings of KDD*.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. 2021. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification. *arXiv preprint arXiv:2108.02035*.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. K-bert: Enabling language representation with knowledge graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2901–2908.
- Robert L Logan IV, Nelson F Liu, Matthew E Peters, Matt Gardner, and Sameer Singh. 2019. Barack’s wife hillary: Using knowledge-graphs for fact-aware language modeling. *arXiv preprint arXiv:1906.07241*.
- Ndapandula Nakashole and Tom M. Mitchell. 2014. **Language-aware truth assessment of fact candidates**. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1009–1019, Baltimore, Maryland. Association for Computational Linguistics.
- Allen Newell, J. C. Shaw, and Herbert A. Simon. 1959. Report on a general problem-solving program. In *IFIP Congress*.
- Liangming Pan, Wenhua Chen, Wenhan Xiong, Min-Yen Kan, and William Yang Wang. 2021. **Zero-shot fact verification by claim generation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 476–483, Online. Association for Computational Linguistics.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. **Hypothesis only baselines in natural language inference**. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI ’09*, page 452–461, Arlington, Virginia, USA. AUAI Press.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M. Marlin. 2013. **Relation extraction with matrix factorization and universal schemas**. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 74–84, Atlanta, Georgia. Association for Computational Linguistics.
- Tal Schuster, Darsh Shah, Yun Jie Serene Yeo, Daniel Roberto Filizzola Ortiz, Enrico Santus, and Regina Barzilay. 2019. **Towards debiasing fact verification models**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3419–3425, Hong Kong, China. Association for Computational Linguistics.
- Joe Stacey, Pasquale Minervini, Haim Dubossarsky, and Marek Rei. 2022. **Logical reasoning with span-level predictions for interpretable and robust NLI models**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3809–3823, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer, and Ido Dagan. 2018. **Supervised open information extraction**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 885–895, New Orleans, Louisiana. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018a. Fever: a large-scale dataset for fact extraction and verification. In *Proceedings of NAACL-HLT*.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2018b. The FEVER2.0 shared task. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2019. **The FEVER2.0 shared task**. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 1–6, Hong Kong, China. Association for Computational Linguistics.

Denny Vrandečić and Markus Krötzsch. 2014. [Wikidata: A free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788.

Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021a. Kepler: A unified model for knowledge embedding and pre-trained language representation. *Transactions of the Association for Computational Linguistics*, 9:176–194.

Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021b. [Kepler: A unified model for knowledge embedding and pre-trained language representation](#). *Transactions of the Association for Computational Linguistics*, 9:176–194.

Dustin Wright, David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Isabelle Augenstein, and Lucy Lu Wang. 2022. Generating scientific claims for zero-shot scientific fact checking. *arXiv preprint arXiv:2203.12990*.

Alexander Yates, Michele Banko, Matthew Broadhead, Michael Cafarella, Oren Etzioni, and Stephen Soderland. 2007. [TextRunner: Open information extraction on the web](#). In *Proceedings of Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, pages 25–26, Rochester, New York, USA. Association for Computational Linguistics.

Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2018. Neural question generation from text: A preliminary study. In *Natural Language Processing and Chinese Computing: 6th CCF International Conference, NLPCC 2017, Dalian, China, November 8–12, 2017, Proceedings 6*, pages 662–671. Springer.

## A Hyperparameters

For the universal schema model, the hyperparameters were manually tuned using the wikidata dev set, with a batch size of 32, a learning rate of  $2e-5$ , an Adam epsilon of  $1e-8$ , a weighted decay of 0.01 and a maximum gradient norm of 1.0. The model was trained for a maximum of 3 epochs, with early stopping based on the loss observed on the dev set. Given the large amount of training data and limited computing resources, we partition the data into sections of 10000000 randomly shuffled samples to make the task feasible. Each section is treated as a separate training batch for our model.