

# Zero-Shot Learning and Key Points Are All You Need for Automated Fact-Checking

Mohammad Ghiasvand Mohammadkhani<sup>1</sup>, Ali Ghiasvand Mohammadkhani<sup>2</sup>  
Hamid Beigy<sup>3</sup>

<sup>1</sup>Amirkabir University of Technology, <sup>2</sup>Shahid Soltani 4 High School

<sup>3</sup>Sharif University of Technology

mohammad.ghiasvand@aut.ac.ir, aghiasvandm@gmail.com, beigy@sharif.edu

## Abstract

Automated fact-checking is an important task because determining the accurate status of a proposed claim within the vast amount of information available online is a critical challenge. This challenge requires robust evaluation to prevent the spread of false information. Modern large language models (LLMs) have demonstrated high capability in performing a diverse range of Natural Language Processing (NLP) tasks. By utilizing proper prompting strategies, their versatility—due to their understanding of large context sizes and zero-shot learning ability—enables them to simulate human problem-solving intuition and move towards being an alternative to humans for solving problems. In this work, we introduce a straightforward framework based on *Zero-Shot Learning* and *Key Points* (ZSL-KeP) for automated fact-checking, which despite its simplicity, performed well on the AVeriTeC shared task dataset by robustly improving the baseline and achieving 10<sup>th</sup> place.<sup>1</sup>

## 1 Introduction

The AVeriTeC task (Schlichtkrull et al., 2024) is designed to encourage the development of advanced frameworks for automated fact-checking, a critical task in NLP. With the rapid spread of information and misinformation online, automated fact-checking is increasingly important. Given the time-consuming nature of manual fact-checking, building an effective neural language model-based framework is valuable for saving time and costs, improving performance, and supporting human judgment. Significant efforts are being made to automate this process within digital tools or LLMs (Nakov et al., 2021).

LLMs with billions of parameters offer extensive knowledge and strong reasoning capabilities that

can be customized for various tasks. Designing effective and appropriate prompts is crucial in this customization process. Recent utilization of LLMs can mainly be divided into two categories: fine-tuning and In-Context Learning (ICL). Given the enormous size of LLMs and the high computational cost associated with fine-tuning them, utilizing ICL through zero-shot or few-shot prompting is much more efficient.

Explaining the reasoning behind a decision is crucial for user trust in automated fact-checking, as users need to understand the evidence behind the model’s verdict (Guo et al., 2022). This work employs Large Language Models (LLMs) with Zero-Shot Learning (ZSL), which offer advantages over simpler, classification-based models due to their long context windows and high reasoning capabilities. Besides using powerful LLMs and effective prompting, accurate retrieval of relevant information is vital. This involves hierarchical, step-by-step prompting and decomposition-based retrieval methods (Zhang and Gao, 2023). This paper describes the novel approach implemented by our team, **MA-Bros-H**, for the AVeriTeC shared task, which integrates ZSL and key point utilization within a unified and straightforward framework.

## 2 Related Works

To highlight a few recent research efforts in automated fact-checking, it is notable that (Kotonya and Toni, 2020) provided explainability through summarization, and (Lee et al., 2020) utilized the internal knowledge of pretrained language models such as BERT (Devlin, 2018) within their framework. Additionally, (Lee et al., 2021) employed few-shot prompting for fact-checking, (Zhang and Gao, 2023) introduced a hierarchical, step-by-step prompting method that involves claim decomposition followed by step-by-step reasoning to predict the final verdict, and (Kim et al., 2024) proposed

<sup>1</sup>Code and data released at <https://github.com/mghiasvand1/ZSL-KeP>




Figure 1: ZSL-KeP Framework Illustration

a multi-agent debate strategy for explainable fact-checking.

### 3 Methodology

This section provides a detailed overview of the problem definition of AVeriTeC task, as well as the operational procedure of our ZSL-KeP model, as outlined in Figure 1.

#### 3.1 Problem Definition

Since our method is explicitly based on zero-shot prompting, we use only the test data to execute our framework, ignoring the train and validation datasets. For each data point in the test dataset, a claim is provided, and a verdict must be predicted from the labels “*Supported*”, “*Refuted*”, “*Not Enough Evidence*”, and “*Conflicting Evidence/Cherry-Picking*”. Additionally, for each claim, a JSON file called a knowledge store is provided. This file contains numerous URLs with scraped texts, including some gold documents that assist in selecting the accurate label. The expected output includes a verdict for the input claim and adequate, yet non-redundant, evidence, preferably in the form of question-answer pairs, along with the corresponding URL and scraped text for each pair to justify the source of each proposed question-answer pair. It is noteworthy that the answer type for each question can be “*Extractive*”, “*Abstractive*”, “*Boolean*” or “*Unanswerable*”.

#### 3.2 ZSL-KeP Framework

Our ZSL-KeP framework is a procedure that contains multiple steps detailed below. However, compared to the baseline method proposed in (Schlichtkrull et al., 2024), our method is much more straightforward, containing fewer steps than the baseline, does not require any fine-tuning, and is simpler to implement.

##### 3.2.1 Zero-Shot Key Points Construction

In the first step, we receive the claim as input and aim to construct key points based on the received claim using ZSL with our chosen LLM. The primary objective of forming key points is that even a simple claim can contain several key points. When searching and retrieving information from the knowledge store, more extensive retrieval typically yields more comprehensive information. A claim might not return many helpful documents when queried directly, but by constructing diverse key points from it, we can obtain more relevant and diverse information. As shown in the prompting template in Appendix A, we limit the number of primitive key points to four. For these distinct key points, we ask the LLM to identify and return pairs of key points whose combinations result in valuable and richer key points. This process aims to construct an extensive set of key points based on the input claim, facilitating more divergent retrieval in the next step.

##### 3.2.2 Extensive Retrieval with References

As mentioned, for each claim, we have a large knowledge store consisting of various URLs with their scraped texts, among which the gold documents for selecting the best and correct verdict are present. In the previous step, we constructed several key points for each claim, either of a normal type or paired, as explained earlier. If the number of constructed key points is  $n$ , we treat these key points as a list of queries. We append the main input claim to this list and use BM25 (Robertson et al., 2009) to retrieve results for each of the  $n + 1$  queries with a different  $top\_k$  parameter for each query. For each selected retrieval result, since each JSON file contains many URLs and each URL has several scraped texts, we construct an ID by concatenating the URL index within the JSON file

Method	Q only	Q + A	AVeriTeC score
AVeriTeC Baseline (Schlichtkrull et al., 2024)	0.24	0.20	0.11
ZSL-KeP (Ours)	0.38	0.24	0.27

Table 1: Main results include retrieval scores for both questions alone and for questions with answers, as well as the AVeriTeC score for the baseline and our proposed method.

with an underscore, followed by the index of the scraped text within the list. For each retrieval document, we attach the text “<ID>” (where ID is the constructed corresponding ID) to the document. After retrieving and appending all these documents for each query, we separate them with a newline character. Finally, we concatenate all groups of retrievals, separating them with two newline characters and several dashes in between, to form a unified retrieval string for the input claim.

### 3.2.3 Zero-Shot Prediction

In this stage, which is the final step of our framework, we use ZSL to generate evidence, followed by a justification and, finally, a verdict. We pass the original claim along with the unified retrieval string formed in the previous step as input, exactly as shown in Appendix A; However, due to the limited context window of the LLM we are using, errors may arise. In such cases, we reduce the number of documents in the unified retrieval string and prompt the LLM again with a shorter input length. The reason we include only the retrievals in the unified retrieval string and omit the key points is that we want to avoid influencing the evidence construction process—specifically, the creation of question-answer pairs—in our strategy. We aim to keep this process dynamic, based on the available selected knowledge and the claim’s purpose.

Since the number of adequate question-answer pairs available as evidence for any claim may vary, we limit the LLM to providing at most 4 pairs to avoid penalties from additional, non-essential question-answer pairs in our prompt. The justification is needed to reason about the verdict based on the evidence and to directly write the predicted verdict afterward. Since the task requires the URL and scraped text for each item of evidence, we instruct the LLM to provide the citation ID when answering questions. This ensures that we can show the source for our verdict and each question-answer pair.

## 4 Experiments and Results

### 4.1 Experimental Setup

In this work, we utilized the *Llama-3-70B model* for both steps described in Sections 3.2.1 and 3.2.3, using the Groq API<sup>2</sup>. Additionally, we set the *temperature* to 0 to ensure reproducibility and *top\_p* to 0.8. For key point construction, we set *max\_length* to 512, and for zero-shot prediction, we set it to 1024. In the retrieval step using BM25, we set *top\_k* to 70 for the original claim and to 12 for other queries, which include key points from both normal and combined forms. For zero-shot prediction, which is the third step of the strategy, if a rate limit occurs due to input length limitations, we retain only the first 55 documents for the original claim and 9 documents for key point retrievals.

### 4.2 Evaluation Metrics

The AVeriTeC scoring follows a similar approach to FEVER (Thorne et al., 2018) and considers the correctness of the verdict label conditioned on the correctness of the evidence retrieved. The label will only be considered correct if it matches with the gold label and the Hungarian meteor score between the predicted evidence and the gold evidence is at least 0.25. However, Unlike in FEVER using a closed source of evidence such as Wikipedia, AVERITEC is intended for use with evidence retrieved from the open web. Since the same evidence may be found in different sources, we cannot rely on exact matching to score retrieved evidence. As such, the shared task evaluation strategy instead rely on approximate matching. Specifically, the Hungarian Algorithm (Kuhn, 1955) is used to find an optimal matching of provided evidence to annotated evidence.

### 4.3 Main Results

Despite our framework’s straightforward procedure, which does not require any fine-tuning and only utilizes ZSL, as depicted in Table 1, it robustly improves the baseline in both retrieval

<sup>2</sup><https://groq.com/>

scores—calculated for questions alone and for questions with answers—and the AVeriTeC score. This includes improvements of 0.14, 0.04, and 0.16 in retrieval scores for questions only, retrieval scores for questions with answers, and the AVeriTeC score, respectively. Based on these results, by using an open-source LLM, our framework has achieved a 10<sup>th</sup> rank among all 23 system result submissions.

## 5 Conclusion

In this paper, we introduced ZSL-KeP, an effective yet straightforward framework for automated fact-checking. We utilized the ZSL capability of LLMs and constructed key points for extensive retrieval to generate evidence in a question-and-answer pairs format, along with a final verdict. By relying solely on the ICL capability of LLMs, our strategy operates without requiring any fine-tuning and is more straightforward compared to the baseline. Our framework sets a new benchmark, indicating promising avenues for future research in related topics.

## 6 Limitations

While our work shows strong performance, it has some limitations that suggest areas for future research. Our method improves diversity by using zero-shot key points for retrieval, but the limited input length of our LLM, constrained by time and budget limitations, prevented us from retrieving a larger document set. Additionally, a more powerful LLM could enhance accuracy in generating evidence and verdicts. Addressing these issues could significantly improve our framework’s results.

## References

Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.

Kyungha Kim, Sangyun Lee, Kung-Hsiang Huang, Hou Pong Chan, Manling Li, and Heng Ji. 2024. Can llms produce faithful explanations for fact-checking? towards faithful explainable fact-checking via multi-agent debate. *arXiv preprint arXiv:2402.07401*.

Neema Kotonya and Francesca Toni. 2020. Explainable automated fact-checking for public health claims. In

*Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7740–7754.

Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.

Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. 2021. Towards few-shot fact-checking via perplexity. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1971–1981.

Nayeon Lee, Belinda Z Li, Sinong Wang, Wen-Tau Yih, Hao Ma, and Madian Khabsa. 2020. Language models as fact checkers? *ACL 2020*, page 36.

P Nakov, D Corney, M Hasanain, F Alam, T Elsayed, A Barron-Cedeno, P Papotti, S Shaar, G Da San Martino, et al. 2021. Automated fact-checking for assisting human fact-checkers. In *IJCAI*, pages 4551–4558. International Joint Conferences on Artificial Intelligence.

Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.

Michael Schlichtkrull, Zhijiang Guo, and Andreas Vlachos. 2024. Averitec: A dataset for real-world claim verification with evidence from the web. *Advances in Neural Information Processing Systems*, 36.

James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819.

Xuan Zhang and Wei Gao. 2023. Towards llm-based fact verification on news claims with a hierarchical step-by-step prompting method. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 996–1011.

## A Prompt Templates

This section provides all of the prompting templates used within the strategy. Figure 2 illustrates the full prompts for section 3.2.1, while the prompts for section 3.2.3 are shown in Figure 3. It is noteworthy that in the user messages, the tags “<claim>” and “<retrieval>” are replaced by the original claim and the unified retrieval string, respectively.

[System]  
 You're a helpful assistant with expertise in understanding the concepts of a given claim and writing subclaims that decompose the main claim well.

[User]  
 Based on the given claim, your task is to extract several distinct key points (2 to 4, depending on the claim's length and complexity) without paraphrasing, in the format of short sentences. Focus on key points that support the main intent of the claim, rather than unnecessary details. Then, only if the number of key points is more than two, identify the pairs of key points whose combination leads to new and richer key points, and return a single coherent short text as a representation of each combination without paraphrasing. Provide your response explicitly in the format of {"key\_points": [], "combined\_key\_points": []}.

Claim: <claim>

Figure 2: The Prompts for Zero-Shot Key Points Construction

[System]  
 You are a helpful assistant with expertise in creating evidence through suitable question-answer pairs based on a given claim and the available key points within the retrieved knowledge, and in providing an accurate verdict for that claim.

[User]  
 Your task is to accurately determine a correct verdict for a given claim from the labels "Refuted", "Supported", "Not Enough Evidence", or "Conflicting Evidence/Cherry-Picking". You need to provide 1 to 4 necessary and helpful question-answer (QA) pairs. Each QA pair should be well-constructed, focusing on different important parts of the claim and utilizing the retrieved knowledge effectively to guide accurate decision-making. Therefore, you need to break down the claim into its distinct and most important subclaims, focusing on these individual components, as well as considering direct questions related to the main claim if the retrieved knowledge is sufficient. Your answers can only be in the forms of extractive (preferred), abstractive, or unanswerable. Extractive answers are those directly pulled from the text, while abstractive answers summarize or infer information based on the text. Unanswerable type is very rare, and in this case, set the answer to "No answer could be found." and the citation\_id to "". Each piece of text in the retrieved knowledge has a <citation\_id> at its end, where the placeholder is replaced by the main citation ID. For each proposed answer to all answerable questions in your evidence, you must include exactly one citation ID (if there are multiple citation\_id, select only one) solely within the "citation\_id" field. After providing evidence, you must also provide a concise justification explaining how the evidence and the retrieved knowledge support the selected label for the claim. Provide your answer explicitly in the following format without any other change or additional feedback:

```
{
  "evidence": [
    {
      "question": "question",
      "answer": "answer",
      "citation_id": "<citation_id>"
    },
    ...
  ],
  "justification": "justification",
  "pred_label": "pred_label"
}
```

Claim:  
 <claim>

Retrieved Knowledge:  
 <retrieval>

Figure 3: The Prompts for Zero-Shot Prediction