

# Evidence-backed Fact Checking using RAG and Few-Shot In-Context Learning with LLMs

Ronit Singhal<sup>1</sup>, Pransh Patwa<sup>2</sup>, Parth Patwa<sup>3</sup>,  
Aman Chadha<sup>4,5\*</sup>, Amitava Das<sup>6</sup>,

<sup>1</sup>IIT Kharagpur, India, <sup>2</sup>Aditya English Medium School, India, <sup>3</sup>UCLA, USA,  
<sup>4</sup>Stanford University, USA, <sup>5</sup>Amazon GenAI, USA, <sup>6</sup>University of South Carolina, USA  
<sup>1</sup>ronit@kgpian.iitkgp.ac.in, <sup>2</sup>pransh.patwa@aemspune.edu.in, <sup>3</sup>parthpatwa@g.ucla.edu  
<sup>4,5</sup>hi@aman.ai, <sup>6</sup>amitava@mailbox.sc.edu

## Abstract

Given the widespread dissemination of misinformation on social media, implementing fact-checking mechanisms for online claims is essential. Manually verifying every claim is very challenging, underscoring the need for an automated fact-checking system. This paper presents our system designed to address this issue. We utilize the Averitec dataset (Schlichtkrull et al., 2023) to assess the performance of our fact-checking system. In addition to veracity prediction, our system provides supporting evidence, which is extracted from the dataset. We develop a Retrieve and Generate (RAG) pipeline to extract relevant evidence sentences from a knowledge base, which are then inputted along with the claim into a large language model (LLM) for classification. We also evaluate the few-shot In-Context Learning (ICL) capabilities of multiple LLMs. Our system achieves an 'Averitec' score of 0.33, which is a 22% absolute improvement over the baseline. Our Code is publicly available on <https://github.com/ronit-singhal/evidence-backed-fact-checking-using-rag-and-few-shot-in-context-learning-with-llms>.

## 1 Introduction

The proliferation of fake news and misinformation on social media platforms has emerged as a significant contemporary issue (Panke, 2020). False online claims have, in some cases, incited riots (Lindsay and Grewar, 2024) and even resulted in loss of life (Kachari, 2018). This problem is particularly amplified during critical events such as elections (Bovet and Makse, 2019) and pandemics (Karimi and Gambrell, 2020; Bae et al., 2022; Morales et al., 2021). Given the vast volume of online content, manually fact-checking every claim is impractical. Therefore, the development of an automated fact




Figure 1: Overview diagram of our system. First, the claim is used to retrieve the top 3 relevant documents. Next, evidence is extracted from these documents using questions and answers generated by an LLM. Finally, the evidence is used for veracity prediction via few-shot ICL.

verification system is imperative. Moreover, simply assigning a veracity label is inadequate; the prediction must be supported by evidence to ensure the system’s transparency and to bolster public trust. Although recent solutions have been proposed (Patwa et al., 2021a; Capuano et al., 2023), the problem remains far from resolved and requires further research efforts.

In this paper, we present our system for automated fact verification. Our system classifies a given textual claim into one of four categories: Supported, Refuted, Conflicting Evidence/Cherry-picking, or Not Enough Evidence. Additionally, it provides supporting evidence for the classification. Our approach leverages recent advancements in Large Language Models (LLMs), specifically Retrieval-Augmented Generation (RAG) and In-Context Learning (ICL), to produce evidence-backed veracity predictions. Given a claim and a collection of documents, our system first employs a RAG pipeline to retrieve the three most relevant documents and extract evidence from them. Subsequently, we utilize ICL to determine the veracity of the claim based on the extracted evidence. Figure 1 provides a high-level overview of our system. We evaluate our system on the Averitec dataset (Schlichtkrull et al., 2023), where it outper-

\*Work does not relate to position at Amazon.

forms the official baseline by a large margin. Our key contributions are as follows:

- We develop a system for automated fact verification that integrates RAG with ICL to provide evidence-based classifications.
- Our proposed system requires only a minimal number of training samples, thereby eliminating the need for a large manually annotated dataset.
- We conduct experiments with various recent LLMs and provide a comprehensive analysis of the results.

The remainder of this paper is structured as follows: Section 2 provides a literature review of related works, while Section 3 describes the dataset. In Section 4, we outline our methodology, followed by a detailed account of the experimental setup in Section 5. Section 6 presents and analyzes our results, and finally, we conclude in Section 7.

## 2 Related Work

Recently, there has been increased research interest in fake news detection and fact checking. Glazkova et al. (2021) proposed an ensemble of BERT (Devlin et al., 2019) for Covid fake news (Patwa et al., 2021b) detection. Harrag and Djahli (2022) employed deep learning techniques for fact checking in Arabic (Baly et al., 2018). (Song et al., 2021) tackled the problem of fake news detection using graph neural networks. The factify tasks (Mishra et al., 2022; Suryavardan et al., 2023b) aimed to detect multi-modal fake news. However, these systems only provide the veracity prediction without any evidence.

On the FEVER dataset (Thorne et al., 2018), Krishna et al. (2022) designed a seq2seq model to generate natural logic-based inferences as proofs, resulting in SoTA performance on the dataset. Schuster et al. (2021) released the VitaminC dataset and propose contrastive learning for fact verification. Hu et al. (2022) proposed a DRQA retriever (Chen et al., 2017) based method for fact checking over unstructured information (Aly et al., 2021). These systems provide evidence or explanation to back their predictions but they test the veracity of synthetic claims whereas we test real claims.

Some researchers have also used LLMs to tackle the problem. Kim et al. (2024) leveraged multiple

Class	Train	Dev
Supported	847	122
Refuted	1743	305
Conflicting evidence/Cherrypicking	196	38
Not enough evidence	282	35
<b>Total</b>	<b>3068</b>	<b>500</b>

Table 1: Class-wise distribution of train and dev set of the dataset. The data is skewed towards the Refuted class.




Figure 2: Word cloud of the claims. We can see that Politics and COVID-19 are common topics in the claims.

LLMs as agents to enhance the faithfulness of explanations of evidence for fact-checking. Zhang and Gao (2023) designed a hierarchical prompting method which directs LLMs to separate a claim into several smaller claims and then verify each of them progressively.

There have also been attempts to solve the problem using RAG. Khaliq et al. (2024) utilized multimodal LLMs with a reasoning method called chain of RAG to provide evidence based on text and image. Deng et al. (2024) proposed a method to decrease misinformation in RAG pipelines by re-ranking the documents during retrieval based on a credibility score assigned to them. Similar to these systems, we also use RAG and LLMs in our solution.

For more detailed surveys, please refer to Thorne and Vlachos (2018); Kotonya and Toni (2020); Guo et al. (2022).

## 3 Data

We utilize the Averitec dataset (Schlichtkrull et al., 2023) for fact-checking purposes. This dataset comprises claims accompanied by a knowledge store (a collection of articles). Each claim is annotated with question-answer pairs that represent the evidence, a veracity label, and a justification for the label. The veracity label can be one of

**Convert the following claim to one neutral question. Do not miss out anything important from the claim. Question the claim, not the fact.**

**Claim:** Donald Trump has stated he will not contest for the next elections

**Incorrect Question:** "What did Donald Trump state for the next elections?"

**Correct Question:** "Did Donald Trump state that he will not contest for the next elections?"

**Claim:** [another claim]

**Incorrect Question:** [example of an incorrect question]

**Correct Question:** [expected correct question]

**Given claim:** In a letter to Steve Jobs, Sean Connery refused to appear in an Apple commercial.

**Generated question:** "Is it true that Sean Connery wrote a letter to Steve Jobs refusing to appear in an Apple commercial?"

Figure 3: The prompt used for generating questions. Some manually created correct and incorrect examples are given to guide the LLM.

the following: Support (S), Refute (R), Conflicting Evidence/Cherry-picking (C), or Not Enough Evidence (N). A claim is labeled as C when it contains both supporting and refuting evidence. The data distribution, as shown in Table 1, indicates a class imbalance favoring the R class, while the C and N classes have relatively few examples. The final testing is conducted on 2,215 instances (Schlichtkrull et al., 2024). For further details on the dataset, please refer to Schlichtkrull et al. (2023, 2024).

On average, each claim consists of 17 words. Figure 2 (word cloud of the claims) reveals that most claims are related to politics and COVID-19.

## 4 Methodology

Given a claim and a knowledge store, our system is comprised of three key components: relevant document retrieval, evidence extraction from the documents, and veracity prediction based on the extracted evidence. The first two components form our Retrieval-Augmented Generation (RAG) pipeline.

### 4.1 Document Retrieval Using Dense Embeddings

In the document retrieval phase, it is essential to match claims with relevant documents from

**Your task is to extract a portion of the provided text that directly answers the given question. The extracted information should be a conclusive answer, either affirmative or negative, and concise, without any irrelevant words. You do not need to provide any explanation. Only return the extracted sentence as instructed. You are strictly forbidden from generating any text of your own.**

**Question:** Is it true that Sean Connery wrote a letter to Steve Jobs refusing to appear in an Apple commercial?

**Document text:** [entire text of one of the retrieved documents]

**Generated answer:** "No, it is not true that Sean Connery wrote a letter to Steve Jobs refusing to appear in an Apple commercial. The letter was a fabrication created for a satirical article on Scoopertino."

Figure 4: The prompt used for generating answers. This prompt is repeated for each of the top three documents.

a knowledge store (in our case, the knowledge store consists of documents provided in the dataset, though it could be replaced with documents retrieved via a search engine). To facilitate this, all documents are first transformed into dense vector embeddings using an embedding model. Since our knowledge store is static, this transformation is a one-time process. The claim in question is then converted into embeddings using the same model.

Once the claim is embedded, we utilize FAISS (Facebook AI Similarity Search) (Douze et al., 2024) to conduct a nearest-neighbor search within the knowledge store. FAISS is an efficient library for similarity search and clustering of dense vectors. We configure FAISS to retrieve the top three documents most relevant to the claim. These documents are then used in the subsequent evidence extraction and veracity prediction steps.

### 4.2 Evidence Extraction Using LLMs

After identifying the top three relevant documents, the next step involves extracting evidence supported by these documents. This process consists of two steps:

**Question Generation:** The claim is transformed into a question challenging its validity using an LLM. We employ In-Context Learning, which enables the model to generate responses based on a few provided examples, aiding in the creation of nuanced and contextually appropriate questions.

Classify the given claim based on provided statements into one of:

1. 'Supported' if there is sufficient evidence indicating that the claim is legitimate. 2. 'Refuted' if there is any evidence contradicting the claim.

3. 'Not Enough Evidence' If you cannot find any conclusive factual evidence either supporting or refuting the claim.

4. 'Conflicting Evidence/Cherrypicking' if there is factual evidence both supporting and refuting the claim.

**Claim:** [claim]

**Statements:** [statements related to claim]

**Class:** [ground truth class]

**Claim:** [claim]

**Statements:** [statements related to claim]

**Class:** [ground truth class]

**Given Claim:** New Zealand's new Food Bill bans gardening.

**Given Statements:** ["The Food Bill does not impose restrictions on personal horticultural activities, such as growing vegetables and fruits at home.", "Gardening is not banned in New Zealand.", "There are no laws against people having gardens, or sharing food that they've grown at home, said a spokesperson for New Zealand's Ministry for Primary Industries."]

**Generated class:** Refuted

Figure 5: A prompt similar to the one used for generating the final prediction. The actual prompt has some more instructions which are omitted here in the interest of space. Two annotated train examples are provided for the LLM to learn from.

The prompt is designed to ensure that the generated question challenges the claim's veracity rather than simply seeking a factual answer. An example prompt is provided in Figure 3.

**Answer Generation:** After generating the question, we provide a single document to an LLM and pose the question. The LLM is prompted to deliver concise and definitive answers derived directly from the content of the document. This process is repeated for each of the three documents, resulting in three distinct answers for each claim. These answers collectively constitute our evidence. It is important to note that in our experiments, the LLM used for answer generation does not necessarily need to be the same as the one used for question generation. The prompt utilized in this step is similar to the one depicted in Figure 4.

### 4.3 Few-Shot ICL for Final Classification

For the final veracity prediction, we use an LLM to classify a claim based on the three pieces of evidence extracted earlier. The LLM is prompted

to choose one out of the four possible classes. The prompt is designed to guide the model through the classification process, ensuring that it correctly interprets the relationship between the claim and the evidence. An example prompt is given in Figure 5.

Our methodology aligns with recent advancements in retrieval-augmented generation (RAG) pipelines which alleviate hallucination and ICL methods, which have been shown to improve the accuracy of LLMs. The integration of these state-of-the-art methods is an attempt to ensure that the extracted evidence is both relevant and contextually appropriate for validating the claims accurately.

## 5 Experiments

To convert documents into dense embeddings, we utilize the `dunzhang/stella_en_1.5B_v5` model<sup>1</sup>. This model is chosen because, at the time of our experiments, it was ranked first on the Massive Text Embedding Benchmark (MTEB) leaderboard (Muennighoff et al., 2022), and holds the second position at the time of writing this paper.

For all LLMs used in our experiments, we employ their 4-bit quantized versions via Ollama<sup>2</sup>. This quantization enables us to load larger LLMs onto our GPUs.

For question generation, we use the Phi-3-medium model (Abdin et al., 2024). The temperature is set to 0, and greedy decoding is used to ensure that the answers are as factual as possible and to minimize hallucinations.

For answer generation and final classification, we experiment with multiple LLMs of varying sizes, including InternLM2.5 (Cai et al., 2024), Llama-3.1 (Dubey et al., 2024), Phi-3-medium (Abdin et al., 2024), Qwen2 (Yang et al., 2024), and Mixtral (Jiang et al., 2024). These models are selected based on their performance on the Open LLM Leaderboard (Fourrier et al., 2024) and their availability through Ollama.

We utilize an A40 GPU for Mixtral, while all other models are run on an A100 GPU. Our best-performing model, Mixtral, requires an average of 2 minutes for evidence extraction and final prediction. Our code is publicly available on <https://github.com/ronit-singhal/evidence-backed-fact-checking-using-rag-and-few-shot-in-context->

<sup>1</sup>[https://huggingface.co/dunzhang/stella\\_en\\_1.5B\\_v5](https://huggingface.co/dunzhang/stella_en_1.5B_v5)

<sup>2</sup><https://github.com/ollama/ollama>



Model	Size	Q+A $\uparrow$	Averitec $\uparrow$	Acc $\uparrow$
InternLM2.5	7B	0.278	0.194	0.374
Llama3.1	8B	0.259	0.224	0.538
Phi-3-Medium	14B	0.259	0.28	0.654
Llama 3.1	70B	0.272	0.328	<b>0.662</b>
Qwen2	72B	0.285	0.33	0.61
Mixtral	8*22B	<b>0.292</b>	<b>0.356</b>	0.636

Table 2: Results of various models on the dev set. Performance improves as the model size increases. Acc refers to accuracy. Q+A and Averitec scores are described in Section 5.1.

System	Q $\uparrow$	Q+A $\uparrow$	Averitec $\uparrow$
Official Baseline	0.24	0.2	0.11
Mixtral (ours)	<b>0.35</b>	<b>0.27</b>	<b>0.33</b>

Table 3: Results on the test set. Our system which uses Mixtral for final prediction outperforms the official baseline in all metrics. For more details of the metrics, please refer to section 5.1.

## learning-with-llms.

### 5.1 Evaluation Metrics

The evaluation metrics used ensure that credit for a correct veracity prediction is given only when the correct evidence has been identified.

To evaluate how well the generated questions and answers align with the reference data, the pairwise scoring function METEOR (Banerjee and Lavie, 2005) is used. The Hungarian Algorithm (Kuhn, 1955) is then applied to find the optimal matching between the generated sequences and the reference sequences. This evidence scoring method is referred to as Hungarian METEOR. The system is evaluated on the test set using the following metrics:

- **Q only:** Hungarian METEOR score for the generated questions.
- **Q + A:** Hungarian METEOR score for the concatenation of the generated questions and answers.
- **Averitec Score:** Correct veracity predictions where the Q+A score is greater than or equal to 0.25. Any claim with a lower evidence score receives a score of 0.

## 6 Results and Analysis

Table 2 provides a summary of the performance of various models on the development set. The

Model	S	R	N	C	Macro
Mixtral	0.605	0.780	0.126	0.117	<b>0.47</b>
Qwen2	<b>0.620</b>	0.754	<b>0.157</b>	<b>0.153</b>	0.42
Llama 3.1 70b	0.613	<b>0.809</b>	0.022	0	0.361

Table 4: Class-wise F1 scores of our top three LLMs on the dev set. Classes are Supported (S), Refuted (R), Not enough evidence (N), and conflicting evidence/cherrypicking. Macro-averaged F1 score is also reported.

Mixtral 8\*22B model (Jiang et al., 2024) achieves the highest Averitec score, while the Llama 3.1 model (Dubey et al., 2024) attains the highest accuracy. These findings indicate that model performance generally improves with increasing model size. Moreover, the relative rankings of these models on the development set differ from their positions on the Open LLM leaderboard (Fourrier et al., 2024), suggesting that superior performance on the Open LLM leaderboard does not necessarily correlate with better performance in the fact verification task.

Given that Mixtral achieves the highest Averitec score on the development set, we select it for evaluation on the test set. Table 3 provides a comparison of our system and the official baseline (Schlichtkrull et al., 2023) on the test set. The baseline model utilizes Bloom (Scao et al., 2023) for evidence generation, followed by re-ranking of the evidence using a finetuned BERT-large model and finally a finetuned BERT-large model veracity prediction. Unlike the baseline, which uses finetuned models, we only use a few train examples via ICL. Despite that, our system outperforms the baseline across all three evaluation metrics. Notably, our Averitec score of 0.33 is a 22% absolute improvement over the baseline.

### 6.1 Class-wise Performance

Table 4 presents the class-wise performance of our top three models on the development set. Across all models, the Refuted class emerges as the easiest to predict, while the "Not Enough Evidence" and "Conflicting Evidence/Cherrypicking" classes present greater challenges. Notably, no single model excels across all classes. Although Mixtral achieves the highest macro F1 score, it is not the top-performing model for any individual class. Qwen2 surpasses the other models in performance across all classes except Refuted. This suggests that exploring ensemble techniques could be a valu-




Figure 6: Confusion matrix of Mixtral on the development set, illustrating the model’s performance across four classes (C, N, R, S). While class R is mostly accurately classified, classes C and N are often mis-predicted as R or N.

able direction for future research.

Figure 6 illustrates the confusion matrix of Mixtral 8\*22B on the development set. It reveals that both the N and C classes are equally likely to be misclassified as the R and S classes. Additionally, there is significant confusion between the S and R classes, highlighting the inherent difficulty of fact verification.

## 7 Conclusion and Future Work

In this paper, we introduced our system for evidence-supported automated fact verification. Our system - based on RAG and ICL - requires only a minimal number of training examples to extract relevant evidence and make veracity predictions. We observed that all LLMs demonstrate sub-optimal performance on the "Conflicting Evidence/Cherry-picking" and "Not Enough Evidence" categories, which emphasizes the inherent challenges of these categories. Additionally, no single LLM consistently outperforms others across all categories. Our system achieved an Averitec score of 0.33, highlighting the complexity of the problem and indicating a substantial potential for future improvement.

Future research could involve fine-tuning the LLM using parameter-efficient fine-tuning (PEFT) techniques (Liu et al., 2022; Patwa et al., 2024) and improving performance through the use of ensemble techniques (Mohammed and Kora, 2022). Extending the system to include multi-modal fact

verification (Patwa et al., 2022; Suryavardan et al., 2023a) also represents an interesting direction for further investigation.

## 8 Limitation

As we are using few-shot ICL, our system cannot make use of large annotated datasets if available, because of the limitation of the prompt size. Furthermore, we assume the availability of high-quality LLMs, which might not be the case for some low-resource languages.

## 9 Ethical Statement

LLMs are prone to hallucination. In our case, the extracted evidence could be incorrect due to hallucination. Furthermore, the prompts can be tweaked to intentionally generate wrong evidence or predictions. We caution the reader to be aware of such issues and to not misuse the system.

## References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, et al. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *Preprint*, arXiv:2404.14219.
- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. [The fact extraction and VERification over unstructured and structured information \(FEVEROUS\) shared task](#). In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 1–13, Dominican Republic. Association for Computational Linguistics.
- Joseph Bae, Darshan Gandhi, Jil Kothari, Sheshank Shankar, Jonah Bae, Parth Patwa, Rohan Sukumaran, Aviral Chharia, Sanjay Adhikesaven, Shloak Rathod, Irene Nandutu, Sethuraman TV, Vanessa Yu, Kru-tika Misra, Srinidhi Murali, Aishwarya Saxena, Kasia Jakimowicz, Vivek Sharma, Rohan Iyer, Ashley Mehra, Alex Radunsky, Priyanshi Katiyar, Ananthu James, Jyoti Dalal, Sunaina Anand, Shailesh Advani, Jagjit Dhaliwal, and Ramesh Raskar. 2022. [Challenges of equitable vaccine distribution in the covid-19 pandemic](#). *Preprint*, arXiv:2012.12263.
- Ramy Baly, Mitra Mohtarami, James Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018. Integrating stance detection and fact checking in a unified corpus. *arXiv preprint arXiv:1804.08012*.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Ex-*

- trinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Alexandre Bovet and Hernán A. Makse. 2019. [Influence of fake news in twitter during the 2016 us presidential election](#). *Nature Communications*, 10(1):7.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, et al. 2024. [Internlm2 technical report](#). *Preprint*, arXiv:2403.17297.
- Nicola Capuano, Giuseppe Fenza, Vincenzo Loia, and Francesco David Nota. 2023. Content-based fake news detection with machine and deep learning: a systematic review. *Neurocomputing*, 530:91–103.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading wikipedia to answer open-domain questions](#). *Preprint*, arXiv:1704.00051.
- Boyi Deng, Wenjie Wang, Fengbin Zhu, Qifan Wang, and Fuli Feng. 2024. [Cram: Credibility-aware attention modification in llms for combating misinformation in rag](#). *Preprint*, arXiv:2406.11497.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. [The faiss library](#). *Preprint*, arXiv:2401.08281.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, et al. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Clémentine Fourrier, Nathan Habib, Alina Lozovskaya, Konrad Szafer, and Thomas Wolf. 2024. [Open llm leaderboard v2](#). [https://huggingface.co/spaces/open-llm-leaderboard/open\\_llm\\_leaderboard](https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard).
- Anna Glazkova, Maksim Glazkov, and Timofey Trifonov. 2021. [g2tmn at Constraint@AAAI2021: Exploiting CT-BERT and Ensembling Learning for COVID-19 Fake News Detection](#), page 116–127. Springer International Publishing.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Fouzi Harrag and Mohamed Khalil Djahli. 2022. [Arabic fake news detection: A fact checking based deep learning approach](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 21(4).
- Nan Hu, Zirui Wu, Yuxuan Lai, Xiao Liu, and Yansong Feng. 2022. [Dual-channel evidence fusion for fact verification over texts and tables](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5232–5242, Seattle, United States. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, et al. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Panjuri Kachari. 2018. [Death by 'fake news': social media-fuelled lynchings shock india](#). *France24*.
- Nasser Karimi and Jon Gambrell. 2020. [Hundreds die of poisoning in iran as fake news suggests methanol cure for virus](#). *Times of Israel*.
- M Abdul Khaliq, P Chang, M Ma, Bernhard Pflugfelder, and F Miletic. 2024. [Ragar, your falsehood radar: Rag-augmented reasoning for political fact-checking using multimodal large language models](#). *arXiv preprint arXiv:2404.12065*.
- 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](#). *Preprint*, arXiv:2402.07401.
- Neema Kotonya and Francesca Toni. 2020. Explainable automated fact-checking: A survey. *arXiv preprint arXiv:2011.03870*.
- Amrith Krishna, Sebastian Riedel, and Andreas Vlachos. 2022. [ProofVer: Natural logic theorem proving for fact verification](#). *Transactions of the Association for Computational Linguistics*, 10:1013–1030.
- H. W. Kuhn. 1955. [The hungarian method for the assignment problem](#). *Naval Research Logistics Quarterly*, 2(1-2):83–97.
- Martin Lindsay and Calum Grewar. 2024. [Social media misinformation 'fanned riot flames'](#). *BBC*.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohhta, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022. [Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning](#). *Preprint*, arXiv:2205.05638.
- Shreyash Mishra, S Suryavardan, Amrit Bhaskar, Parul Chopra, Aishwarya N Reganti, Parth Patwa, Amitava Das, Tanmoy Chakraborty, Amit P Sheth, Asif Ekbal, et al. 2022. [Factify: A multi-modal fact verification dataset](#). In *DE-FACTIFY@ AAI*.
- Ammar Mohammed and Rania Kora. 2022. An effective ensemble deep learning framework for text classification. *Journal of King Saud University-Computer and Information Sciences*, 34(10):8825–8837.

- Manuel Morales, Rachel Barbar, Darshan Gandhi, Sanskruti Landage, Joseph Bae, Arpita Vats, Jil Kothari, Sheshank Shankar, Rohan Sukumaran, Himi Mathur, Krutika Misra, Aishwarya Saxena, Parth Patwa, Sethuraman T. V., Maurizio Arseni, Shailesh Advani, Kasia Jakimowicz, Sunaina Anand, Priyanshi Katiyar, Ashley Mehra, Rohan Iyer, Srinidhi Murali, Aryan Mahindra, Mikhail Dmitrienko, Saurish Srivastava, Ananya Gangavarapu, Steve Penrod, Vivek Sharma, Abhishek Singh, and Ramesh Raskar. 2021. [Covid-19 tests gone rogue: Privacy, efficacy, mismanagement and misunderstandings](#). *Preprint*, arXiv:2101.01693.
- Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. [Mteb: Massive text embedding benchmark](#). *arXiv preprint arXiv:2210.07316*.
- Stefanie Panke. 2020. [Social media and fake news](#). *aace*.
- Parth Patwa, Mohit Bhardwaj, Vineeth Guptha, Gitanjali Kumari, Shivam Sharma, Srinivas PYKL, Amitava Das, Asif Ekbal, Md Shad Akhtar, and Tanmoy Chakraborty. 2021a. Overview of constraint 2021 shared tasks: Detecting english covid-19 fake news and hindi hostile posts. In *Combating Online Hostile Posts in Regional Languages during Emergency Situation*, pages 42–53, Cham. Springer International Publishing.
- Parth Patwa, Simone Filice, Zhiyu Chen, Giuseppe Castellucci, Oleg Rokhlenko, and Shervin Malmasi. 2024. [Enhancing low-resource LLMs classification with PEFT and synthetic data](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6017–6023, Torino, Italia. ELRA and ICCL.
- Parth Patwa, Shreyash Mishra, S Suryavardan, Amrit Bhaskar, Parul Chopra, Aishwarya Reganti, Amitava Das, Tanmoy Chakraborty, Amit Sheth, Asif Ekbal, and Chaitanya Ahuja. 2022. Benchmarking multimodal entailment for fact verification. In *Proceedings of De-Factify: Workshop on Multimodal Fact Checking and Hate Speech Detection*. CEUR.
- Parth Patwa, Shivam Sharma, Srinivas Pykl, Vineeth Guptha, Gitanjali Kumari, Md Shad Akhtar, Asif Ekbal, Amitava Das, and Tanmoy Chakraborty. 2021b. Fighting an infodemic: Covid-19 fake news dataset. In *Combating Online Hostile Posts in Regional Languages during Emergency Situation*, pages 21–29, Cham. Springer International Publishing.
- Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, et al. 2023. [Bloom: A 176b-parameter open-access multilingual language model](#). *Preprint*, arXiv:2211.05100.
- Michael Schlichtkrull, Yulong Chen, Chenxi Whitehouse, Zhenyun Deng, Mubashara Akhtar, Rami Aly, Zhijiang Guo, Christos Christodoulopoulos, Oana Cocarascu, Arpit Mittal, James Thorne, and Andreas Vlachos. 2024. The automated verification of textual claims (averitec) shared task. In *Proceedings of the Seventh Workshop on Fact Extraction and VERification (FEVER)*.
- Michael Sejr Schlichtkrull, Zhijiang Guo, and Andreas Vlachos. 2023. [Averitec: A dataset for real-world claim verification with evidence from the web](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. [Get your vitamin C! robust fact verification with contrastive evidence](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online. Association for Computational Linguistics.
- Chenguang Song, Kai Shu, and Bin Wu. 2021. [Temporally evolving graph neural network for fake news detection](#). *Information Processing & Management*, 58(6):102712.
- S Suryavardan, Shreyash Mishra, Megha Chakraborty, Parth Patwa, Anku Rani, Aman Chadha, Aishwarya Reganti, Amitava Das, Amit Sheth, Manoj Chinakotla, Asif Ekbal, and Srijan Kumar. 2023a. [Findings of factify 2: Multimodal fake news detection](#). *Preprint*, arXiv:2307.10475.
- S Suryavardan, Shreyash Mishra, Parth Patwa, Megha Chakraborty, Anku Rani, Aishwarya Reganti, Aman Chadha, Amitava Das, Amit Sheth, Manoj Chinakotla, et al. 2023b. [Factify 2: A multimodal fake news and satire news dataset](#). *arXiv preprint arXiv:2304.03897*.
- James Thorne and Andreas Vlachos. 2018. Automated fact checking: Task formulations, methods and future directions. *arXiv preprint arXiv:1806.07687*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, et al. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Xuan Zhang and Wei Gao. 2023. [Towards llm-based fact verification on news claims with a hierarchical step-by-step prompting method](#). *Preprint*, arXiv:2310.00305.