

Loki: An Open-Source Tool for Fact Verification

Haonan Li^{1,2} Xudong Han^{1,2} Hao Wang¹,
Yuxia Wang^{1,2} Minghan Wang³ Rui Xing^{2,4} Yilin Geng^{1,4}
Zenan Zhai¹ Preslav Nakov² Timothy Baldwin^{1,2,4}

¹LibrAI ²MBZUAI ³Monash University ⁴The University of Melbourne

Abstract

We introduce LOKI, an open-source tool designed to address the growing problem of misinformation. LOKI adopts a human-centered approach, striking a balance between the quality of fact-checking and the cost of human involvement. It decomposes the fact-checking task into a five-step pipeline: breaking down long texts into individual claims, assessing their check-worthiness, generating queries, retrieving evidence, and verifying the claims. Instead of fully automating the claim verification process, LOKI provides essential information at each step to assist human judgment, especially for general users such as journalists and content moderators. Moreover, it has been optimized for latency, robustness, and cost efficiency at a commercially usable level. LOKI is released under an MIT license and is available on GitHub.¹ We also provide a video presenting the system and its capabilities.²

1 Introduction

In today’s digital landscape, the rapid spread of misinformation has become a significant societal problem, with far-reaching consequences for politics, public health, and social stability (Pan et al., 2023; Augenstein et al., 2024). With the rise of online platforms, users are exposed to large volumes of information, often without the ability to assess its accuracy. While manual fact-checking is reliable, it is labor-intensive, time-consuming, and often requires domain expertise, creating a gap where misinformation can spread unchecked and cause harm before being addressed.

To address this problem, automated fact-checking systems have been proposed, but have mostly focused on full automation, which can negatively impact quality.

Here, we propose LOKI, which offers a semi-automated, human-in-the-loop approach to fact verification. Instead of completely eliminating human participation, LOKI assists users by breaking down the fact-checking process into five manageable steps, ensuring that human judgment remains integral to decision-making. This benefits users, such as journalists and content moderators, who need reliable tools to quickly and accurately verify information.

LOKI offers a five-step pipeline for fact verification: identifying claims, assessing their check-worthiness, generating queries for evidence retrieval, retrieving evidence, and verifying the claims. This modular framework ensures flexibility and adaptability. It supports fact-checking in multiple languages and integration with large language models (LLMs). Additionally, LOKI is optimized for practical use, with improvements in latency, robustness, and cost-efficiency, making it well-suited for commercial applications.

LOKI is implemented in Python and offers ease of use through multiple interfaces: a user-friendly graphical interface, command-line functionality, and the ability to be imported as a package into other projects. To demonstrate its utility, we compare LOKI against several recent fact-checking tools, and show that optimized system implementation and prompt engineering yields demonstrable improvements over other tools in terms of efficiency while achieving very competitive performance.

2 Related Work

Numerous automated fact-checking systems have been developed, including *RARR*, *FActScore*, *FactTool*, *Factcheck-GPT*, and *Longform SAFE* (Gao et al., 2023; Min et al., 2023; Chern et al., 2023; Wang et al., 2024a; Wei et al., 2024; Fadeeva et al., 2024). However, these tools are often inaccessi-

¹<https://github.com/Libr-AI/OpenFactVerification>

²https://www.youtube.com/watch?v=L_3Dp41Lk_k

Fact-checking System	UI	Asynchronous	Multilingual	Multi-LLM	Granularity	Transparency	Open-source
<i>RARR</i> (Gao et al., 2023)	✗	✗	✗	✗	Document	✗	✓
<i>FActScore</i> (Min et al., 2023)	✗	✗	✗	✗	Claim	✗	✓
<i>FactTool</i> (Chern et al., 2023)	✗	✓	✗	✗	Claim	✗	✓
<i>Factcheck-GPT</i> (Wang et al., 2024a)	✗	✗	✗	✗	Claim	✓	✓
<i>CoVe</i> (Dhuliawala et al., 2024)	✗	✗	✗	✗	Claim	✗	✓
<i>Longform SAFE</i> (Wei et al., 2024)	✗	✗	✗	✓	Claim	✗	✓
<i>FIRE</i> (Xie et al., 2024)	✗	✗	✗	✗	Claim	✓	✓
<i>Perplexity.ai</i>	✓	✓	✓	unclear	Claim	✗	✗
<i>OPENFACTCHECK</i> (Iqbal et al., 2024)	✓	✓	✗	✗	Claim	✓	✓
LOKI (ours)	✓	✓	✓	✓	Claim	✓	✓

Table 1: Comparison of representative automatic fact-checking *pipelines*, DEMOS and products in the last two years from seven perspectives: (1) **UI** — the system has user interface supporting easy interaction with general users; (2) **Asynchronous** processing for retrieving evidence from web pages and calling LLM APIs; (3) **Multilingual** — the system is designed to support languages other than English; (4) **Multi-LLM** — flexibly calling different LLM APIs as fact verifiers; (5) **Granularity** — the smallest granularity of document decomposition and verification supported by the system; (6) **Transparency** — the system can show fine-grained snippets of evidence with the corresponding URL, and the relationship between the evidence and the claim (support, refute, or irrelevant); and (7) **Open-source** — the system is open-sourced.

ble to general users who may not have a Python environment to compile code and run verification processes. Although these systems can serve as backends for various services, they lack a user-friendly web interface that allows users to verify text inputs by simply typing or pasting text and clicking a check button. LOKI addresses this gap by providing an accessible, human-friendly user interface.

Each fact-checking system also has its own strengths. For instance, *Factcheck-GPT* offers a fine-grained framework encompassing all possible subtasks to enhance the fact-checking process. *FactTool* uses a low-latency evidence retriever through asynchronous processing, while *FActScore* introduces a scoring metric that calculates the percentage of true claims within a text, providing a quantitative measure of the input’s credibility. LOKI integrates these advantages into a unified system (Wang et al., 2024b). Table 1 compares eight representative fact-checking pipelines, demos, or products (e.g., *Perplexity.ai*) with LOKI across seven dimensions. Of the systems surveyed, LOKI is the only one to support all seven listed system aspects.

3 Loki Fact Checker

In this section, we provide an in-depth overview of LOKI, focusing on its core functionalities and architectural design. We first outline the five-component fact verification pipeline and then discuss the architecture’s flexibility that allows component substitution for different domains. After that, we explore how parallelism is employed to enhance the sys-

tem’s efficiency and reduce latency. Finally, we present the LOKI user interface, designed to provide a seamless and user-friendly user experience.

3.1 Fact Verification Pipeline

Previous work structures fact verification into a series of steps, commonly including text decomposition, checkworthiness identification, evidence retrieval and collection, stance detection, and correction determination (Wang et al., 2024a). In this work, we propose a five-step pipeline consisting of the following modules: Decomposer, Checkworthiness Identifier, Query Generator, Evidence Retriever, and Claim Verifier.

Decomposer breaks down long texts into smaller, atomic claims. It produces individual claims that can be verified independently. For a better user experience, we ensure the decomposed claims are traceable to the original text. During the result presentation, LOKI displays both the original and decomposed claims for contextual clarity.

Checkworthiness Identifier filters out unworthy claims that are vague, ambiguous, or opinion-based, ensuring only factual statements proceed for verification. For instance, claims like *MBZUAI has a vast campus* are deemed unworthy due to the subjective interpretation of *vast*.

Query Generator converts check-worthy claims into optimized queries for evidence retrieval, focusing on keyword-based retrieval.

Evidence Retriever gathers relevant information to support the verification process. Currently, LOKI retrieves evidence from online sources via search engine APIs.

Claim Verifier evaluates retrieved evidence to verify the claim, presenting supporting or refuting snippets for users to make informed judgments.

3.2 Implementation and Extensibility

LOKI is implemented in Python and can be easily integrated into other projects as an importable package. While the core functionality leverages the capabilities of LLMs, we also provide implementations using traditional NLP tools like NLTK or SpaCy.³ This section introduces LOKI’s core implementation and its extensibility.

LLM-based Implementation LOKI employs LLMs in four of its five components: Decomposer, Checkworthiness Identifier, Query Generator, and Claim Verifier. Each component is implemented as a Python class with functions that interact with LLMs to perform core tasks. For instance, the Decomposer class includes functions to break down long texts into claims and map these claims back to the original text.

Each LLM-based function comes with a set of default prompts, typically including: (1) a brief task description, (2) input and output formats, and (3) few-shot examples to guide the LLM. To generate consistently parsable outputs, we prompt the LLMs to return results in a structured format, such as JSON.⁴ This approach streamlines the integration of LLM outputs into subsequent steps of the pipeline. For each task, we hand-crafted 10 test cases and optimized the prompts to maximize LLM performance. An example prompt for the Decomposer is shown in Figure 1.

To maintain system reliability, all LLM function calls are wrapped with a retry mechanism that automatically retries in case of network issues or transient errors. We use the default hyperparameters for LLM inference and leave hyperparameter tuning to future work.

³Our evaluation shows that LLM-based implementations are generally more robust and accurate in practical scenarios.

⁴JSON output format is currently only well-supported in some models such as GPT4 (OpenAI, 2024). Switching the backend LLM to other models, such as Llama-3, may cause pipeline failures unless specifically designed to ensure a consistent output format.

Your task is to **decompose the text into atomic claims**.
The answer should be a **JSON** with a single key “claims”, with the value of a list of strings, where each string should be a context-independent claim, representing one fact.

Note that:

1. Each claim should be concise (less than 15 words) and self-contained.
2. Avoid vague references like ‘he’, ‘she’, ‘it’, ‘this’, ‘the company’, ‘the man’ and using complete names.
3. Generate at least one claim for each single sentence in the texts.

For example,

Text: *Mary is a five-year old girl, she likes playing piano and she doesn't like cookies.*

Output:

```
{ "claims": [ "Mary is a five-year old girl.", "Mary likes playing piano.", "Mary doesn't like cookies." ] }
```

Text: {doc}

Output:

Figure 1: An example of the prompt for the Decomposer. This is an template and the actual prompt is generated by replacing the placeholder {doc} with the input text.

Evidence Retriever Currently, LOKI uses the Google Search API (Serper API) for open-domain questions, retrieving web-based evidence to support or refute claims.⁵ This module’s design is flexible, allowing integration with alternative search engines or specialized databases to expand evidence sources. For each claim, LOKI generates three distinct search queries.⁶ If the Google Search API returns a direct answer, it is captured as primary evidence. Otherwise, the top five search results are retained, and snippets are extracted. An optional Natural Language Inference (NLI) model can then be applied to rank snippets based on their relevance, prioritizing the most pertinent evidence.⁷

Extensibility LOKI’s modular design enables high extensibility. Each component can be individually optimized or replaced with improved solutions as needed. For LLM-based functions, LOKI supports multiple LLM APIs as well as locally deployed models. Additionally, if users require only specific functionalities, each component can be imported and used independently.

3.3 Parallelism and Efficiency

LOKI uses parallelism to efficiently handle multiple verification requests at the same time. By using Python’s asyncio library, the system can make asynchronous API calls, enabling concurrent processing without blocking the main thread.

⁵<https://serper.dev/>

⁶Using the claim itself as a query can lead to biased or irrelevant results, as false claims may mislead the search.

⁷Snippets classified as either entailment or contradiction are considered highly relevant.

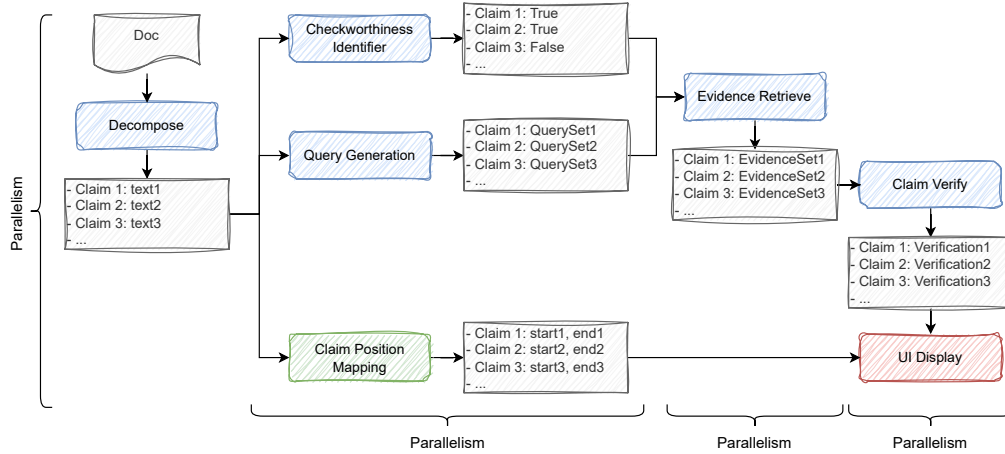


Figure 2: Parallel execution of independent components in the fact-checking pipeline. The entire process can be completed within the combined time of three LLM calls and one web query.

The LLM call function manages these calls and ensures compliance with rate limits with a user-defined maximum request per minute and request window. A traffic queue keeps track of active requests and removes old entries when they exceed the time window. This approach allows LOKI to process high volumes of fact-checking requests efficiently, improving overall system performance and scalability.

We also parallelize all independent components and functions, as shown in Figure 2. This modular design allows different parts of the fact-checking process, such as claim decomposition, query generation, and evidence retrieval, to be executed concurrently. Ideally, the entire fact-checking process can be completed within the time required for three LLM calls and one web query, significantly reducing the overall response time.

3.4 Multilingual Support

LOKI supports multilingual scenarios. Developers can easily adapt the system by modifying prompts for the target language and performing minimal testing. We provide a comprehensive development guide to facilitate this process. Currently, LOKI supports both English and Chinese, with the potential for easy adaptation to other languages.

3.5 Human-in-the-Loop Design

One of the key features of LOKI is its human-in-the-loop design, which prioritizes transparency and user engagement to enhance human decision-making rather than replace it.

Unlike fully automated solutions that aim to

deliver a final verdict without intermediate steps, LOKI presents critical information and insights at each step, to assist users in making well-informed decisions. This approach is especially beneficial for users such as journalists and content moderators, who need reliable tools to verify information while retaining control over the final judgment.

To facilitate user understanding and engagement, LOKI offers information across four distinct levels:

Level 1 — Overall Credibility Score: At the highest level, LOKI provides a single percentage score representing the overall credibility of the input text, based on the proportion of claims verified as *well-supported*, *conflicting*, or *controversial*, offering a quick summary of the text’s reliability.

Level 2 — Claim-Level Analysis: Claims are classified and displayed numerically beside the overall score, with visual cues highlighting each claim directly in the original text, allowing users to quickly identify areas requiring further scrutiny.

Level 3 — Evidence-Level Insight: For each claim, LOKI presents supporting or refuting evidence from various sources. It also provides contextual information to help users understand the broader background, enabling them to make informed judgments.

Level 4 — Detailed Evidence Breakdown: At the finest granularity, LOKI presents detailed information for each piece of evidence, including the source, relevant paragraphs, and the rationale for why the evidence supports or contradicts the claim.

Our interaction design is guided by three core principles: (1) Transparency and Trust — LOKI provides a clear breakdown of information at mul-

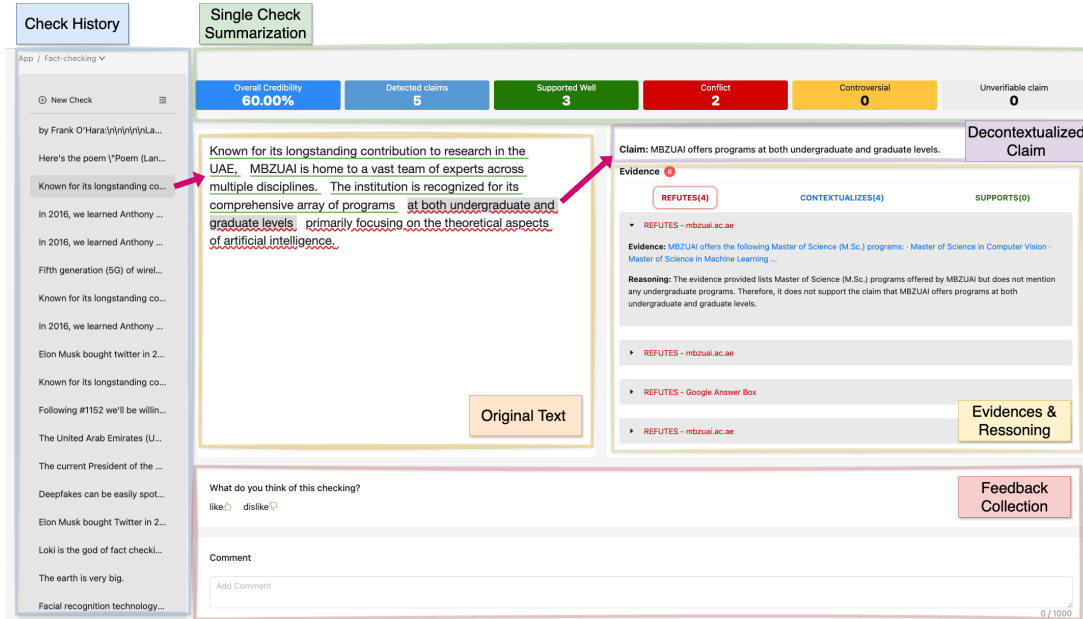


Figure 3: User interface of LOKI. The interface includes sections for Check History (left panel) showing past checks, and Single Check Summarization (top center) summarizing key metrics such as overall credibility and detected claims. Original Text (left center), and the Decontextualized Claim (top right) which isolates the specific claim being evaluated. The Evidence & Reasoning panel (right) provides detailed evidence that refutes, contextualizes, or supports the claim, sourced from various documents, while the Feedback Collection section (bottom) allows users to provide feedback.

multiple levels, allowing users to trace the decision-making process and build trust in the system’s outputs. (2) Hierarchical Presentation — LOKI presents information in layers to avoid overwhelming users, enabling them to “zoom in and out” to different levels of detail for a seamless experience. (3) Assisting, not Replacing: LOKI does not aim to make judgments for the users, but rather to assist them with organized and structured information.

4 System Evaluation and Applications

We performed a comprehensive evaluation for LOKI across several key dimensions, including fact-checking accuracy (measured by precision, recall and F1-score), processing latency (execution wall time), and potential operation cost (number of tokens and queries consumed by API invocation).

Dataset Specifically, our experiments are carried out using two well-known automatic fact-checking evaluation datasets (FacTool-QA (Chern et al., 2023) and Factcheck-Bench (Wang et al., 2024a)), each covering a broad range of domains (see Table 3 for details). In addition to the prompts and LLM responses, these datasets also provide human-annotated claims for each response, along with their labels (True or False), enabling a more

granular evaluation at the claim level rather than at the response level.

Baseline Systems We selected three fact-checking systems and a commercial retrieval-augmented generative model (Perplexity.ai) as baseline systems. While the three fact-checking systems share a similar pipeline, they differ in their components at various stages compared to LOKI, such as the LLM used in the verifier, data sources, and the retriever API being used. To ensure a comprehensive evaluation, we assess the baseline systems across multiple configurations of these components.

Evaluation Protocol To guarantee a fair evaluation across all fact-checking systems, we standardized the evaluation process by bypassing the step of extracting atomic claims from documents. As a result, all systems receive a claim as input and are expected to predict whether the claim is factual. Note that for frameworks like LOKI that predict a factuality score (ranging from 0 to 1), we apply a threshold of 0.8 to convert the score into a binary label.

All systems were evaluated on the same comput-

Framework	Verifier	Source/ Retriever	Factcheck-Bench						FacTool-QA					
			Label = True			Label = False			Label = True			Label = False		
			P	R	F1	P	R	F1	P	R	F1	P	R	F1
Random	–	–	0.79	0.43	0.56	0.18	0.52	0.27	0.79	0.56	0.66	0.28	0.54	0.37
Always True	–	–	0.81	1.00	0.88	0.00	0.00	0.00	0.76	1.00	0.86	0.00	0.00	0.00
Always False	–	–	0.00	0.00	0.00	0.19	1.00	0.33	0.00	0.00	0.00	0.24	1.00	0.39
<i>FactScore</i>	<i>LLaMA</i> 3-Inst 8B	Wiki/BM25	0.87	0.74	0.80	0.34	0.56	0.42	0.82	0.68	0.74	0.34	0.52	0.41
<i>FacTool</i>	<i>LLaMA</i> 3-Inst 8B	Web/Serper	0.88	0.80	0.84	0.40	0.56	0.47	0.93	0.38	0.54	0.32	0.91	0.47
<i>FactScore</i>	GPT-3.5-Turbo	Wiki/BM25	0.87	0.67	0.76	0.31	0.60	0.41	0.82	0.58	0.68	0.31	0.59	0.40
<i>FacTool</i>	GPT-3.5-Turbo	Web/Serper	0.89	0.74	0.81	0.37	0.62	0.46	0.92	0.59	0.72	0.39	0.84	0.53
<i>Factcheck-GPT</i>	<i>GPT-4</i>	Web/SerpAPI	0.90	0.71	0.79	0.52	0.80	0.63	0.88	0.88	0.88	0.63	0.63	0.63
<i>FacTool</i>	<i>GPT-4o</i>	Web/Serper	0.91	0.58	0.71	0.47	0.87	0.61	0.9	0.69	0.78	0.44	0.77	0.56
<i>Perplexity.ai</i>	Sonar-online	Web	0.93	0.73	0.83	0.40	0.76	0.53	0.82	0.88	0.85	0.50	0.38	0.43
LOKI	<i>GPT-4o</i>	Web/Serper	0.84	0.83	0.84	0.63	0.64	0.63	0.89	0.80	0.85	0.53	0.70	0.60

Table 2: The performance of fact-checkers for human-annotated claims in Factcheck-Bench, FacTool-QA, and FELM-WK, judging whether or not a claim is factually true or false with external knowledge (Wikipedia or Web articles) as evidence. *GPT-4* refers to gpt-4-turbo-2024-04-09, *GPT-4o* refers to gpt-4o-2024-05-13.

Dataset ↓	Domain	#True	#False	Total
FacTool-QA	History, geography, biology, science	177	56	233
Factcheck-Bench	Technology, history, science, sports	472	206	678

Table 3: The number of claim labels as well as the involved domain in FacTool-QA, FELM-WK and Factcheck-Bench.

Fact-Checker ↓	Web search (# queries)	# Prompt Tokens	# Completion Tokens	Time (s / sample)
<i>FacTool</i>	3.5 ± 2.3	1895 ± 911	283 ± 212	12.38 ± 46.84
LOKI	4.2 ± 2.6	6518 ± 4055	818 ± 565	8.32 ± 8.17

Table 4: Comparison between *FacTool* vs. LOKI (both are using *GPT-4o* and Serper for verification and retrieval) in terms of cost and latency. Metrics are presented in the form of their average and standard deviation ($\mu \pm \sigma$).

ing node of the M3 Cluster,⁸ with the following configuration: 2 × Intel Xeon Gold 6330 CPUs, 200 GB of memory, and 2 A100 80GB GPUs.

Fact-checking Performance Experimental results for the evaluation of fact-checking accuracy on different systems are presented in Table 2. Although LOKI does not outperform all frameworks on every metric, its overall performance is comparable to the SOTA open-source baseline *Factcheck-GPT*. It’s important to note that Factcheck-GPT uses the more advanced SerpAPI, which retrieves higher-quality documents, significantly improving its performance. However, the cost of SerpAPI (\$0.015 per search) is much higher than Serper (\$0.001 per search), making LOKI the more economical option in terms of cost-effectiveness. When compared to *FacTool*, which uses the same configuration, LOKI outperforms it in most of the metrics across three datasets, demonstrating the

⁸<https://massive.org.au/>

effectiveness of our design.

4.1 Cost and Latency

In Table 4, we provide a detailed comparison of the number of queries, token consumption during LLM API calls, and the average processing time per claim for both LOKI and *FacTool* across the two datasets. We did not include comparisons with other frameworks because only *FacTool* and OPENFACTCHECK support asynchronous processing, which results in significantly lower processing times than the other frameworks.⁹ Since OPENFACTCHECK’s fact-checking pipeline is largely derived from other frameworks, including *FacTool*, we consider *FacTool* sufficiently representative for comparison.

The results show that LOKI consumes more queries and more tokens than *FacTool*, primarily because LOKI retrieves and verifies a greater number of queries and evidence. This also explains why LOKI achieves superior fact-checking performance. Despite handling larger volumes of data, LOKI’s efficient implementation ensures that its average processing time is substantially lower than that of *FacTool*, highlighting its better design, with the asynchronous processing pipeline.

5 Conclusion

In this paper, we present LOKI, a novel open-source Python tool that offers a human-centered approach to fact-checking. LOKI integrates a semi-automated five-step pipeline that balances efficiency with the need for human judgment, particularly for general users like journalists and con-

⁹According to Wang et al. (2025), a cascade pipeline like *Factcheck-GPT* is normally 15.7× slower than an asynchronous pipeline.

tent moderators. Experimental results show that LOKI achieves competitive performance against state-of-the-art systems, offering an optimized and adaptable solution for tackling misinformation. As an open-source tool licensed under MIT, LOKI is accessible for further development, aiming to enhance fact-checking capabilities and support informed decision-making continuously.

References

- Isabelle Augenstein, Timothy Baldwin, Meeyoung Cha, Tanmoy Chakraborty, Giovanni Luca Ciampaglia, David Corney, Renee DiResta, Emilio Ferrara, Scott Hale, Alon Halevy, Eduard Hovy, Heng Ji, Filippo Menczer, Ruben Miguez, Preslav Nakov, Dietram Scheufele, Shivam Sharma, and Giovanni Zagni. 2024. [Factuality challenges in the era of large language models](#). *Nature Machine Intelligence* 6, no. 8, page 852–863.
- I Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, Pengfei Liu, et al. 2023. [Factool: Factuality detection in generative ai—a tool augmented framework for multi-task and multi-domain scenarios](#). *arXiv preprint arXiv:2307.13528*.
- Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2024. [Chain-of-verification reduces hallucination in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3563–3578, Bangkok, Thailand. Association for Computational Linguistics.
- Ekaterina Fadeeva, Aleksandr Rubashevskii, Artem Shelmanov, Sergey Petrakov, Haonan Li, Hamdy Mubarak, Evgenii Tsymbalov, Gleb Kuzmin, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, and Maxim Panov. 2024. [Fact-checking the output of large language models via token-level uncertainty quantification](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 9367–9385, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023. [RARR: Researching and revising what language models say, using language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16477–16508, Toronto, Canada. Association for Computational Linguistics.
- Hasan Iqbal, Yuxia Wang, Minghan Wang, Georgi Nenkov Georgiev, Jiahui Geng, Iryna Gurevych, and Preslav Nakov. 2024. [OpenFactCheck: A unified framework for factuality evaluation of LLMs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 219–229, Miami, Florida, USA. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- OpenAI. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Yikang Pan, Liangming Pan, Wenhui Chen, Preslav Nakov, Min-Yen Kan, and William Wang. 2023. [On the risk of misinformation pollution with large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1389–1403, Singapore. Association for Computational Linguistics.
- Yuxia Wang, Revanth Gangi Reddy, Zain Muhammad Mujahid, Arnav Arora, Aleksandr Rubashevskii, Jiahui Geng, Osama Mohammed Afzal, Liangming Pan, Nadav Borenstein, Aditya Pillai, Isabelle Augenstein, Iryna Gurevych, and Preslav Nakov. 2024a. [Factcheck-bench: Fine-grained evaluation benchmark for automatic fact-checkers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14199–14230, Miami, Florida, USA. Association for Computational Linguistics.
- Yuxia Wang, Minghan Wang, Hasan Iqbal, Georgi Georgiev, Jiahui Geng, and Preslav Nakov. 2025. [Openfactcheck: Build, benchmark customized fact-checking systems and evaluate factuality of claims and llms](#). In *Proceedings of The 31st International Conference on Computational Linguistics (COLING 2025)*, Abu Dhabi, UAE. International Conference on Computational Linguistics.
- Yuxia Wang, Minghan Wang, Muhammad Arslan Manzoor, Fei Liu, Georgi Nenkov Georgiev, Rocktim Jyoti Das, and Preslav Nakov. 2024b. [Factuality of large language models: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19519–19529, Miami, Florida, USA. Association for Computational Linguistics.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#). *CoRR*, abs/2403.18802.
- Zhuohan Xie, Rui Xing, Yuxia Wang, Jiahui Geng, Hasan Iqbal, Dhruv Sahnan, Iryna Gurevych, and Preslav Nakov. 2024. [Fire: Fact-checking with iterative retrieval and verification](#). *arXiv preprint arXiv:2411.00784*.

A Usage of LOKI

Used as a Library LOKI can be integrated as a Python library, allowing developers to incorporate its functionalities directly into their applications. Figure 4 provides an example of how to use FactCheck as a library for verifying text within a Python script.

```
from factcheck import FactCheck

factcheck_instance = FactCheck()

# Example text
text = "Your text here"

# Run the fact-check pipeline
results = factcheck_instance.check_response(text)
print(results)
```

Figure 4: Using FactCheck as a Library

Used as a Web App LOKI can be deployed as a web application, providing a user-friendly interface for interacting with the tool. Figure 5 illustrates the command used to start the web application. The UI of open-source LOKI is introduced in Appendix B

```
python webapp.py --api_config demo_data/api_config.yaml
```

Figure 5: Running FactCheck as a Web App

Multimodal Usage LOKI supports multimodal input, enabling it to process and verify information from various sources, including text, speech, images, and videos. The system can be invoked using different modes specified by the `--modal` argument. Figure 8 demonstrates the usage of FactCheck across these different input modalities

```
# String
python -m factcheck --modal string --input "MBZUAI is the
↪ first AI university in the world"
# Text
python -m factcheck --modal text --input demo_data/text.txt
# Speech
python -m factcheck --modal speech --input
↪ demo_data/speech.mp3
# Image
python -m factcheck --modal image --input
↪ demo_data/image.webp
# Video
python -m factcheck --modal video --input demo_data/video.m4v
```

Figure 8: Multimodal Usage of FactCheck

B User Interface of LOKI

Figure 6 and Figure 7 show the user interface of LOKI (open-source version), which consists of two main pages: the submission page and the result page. The submission page allows users to input text for fact-checking, while the result page displays the decomposed claims, the evidence retrieved for each claim, and the overall credibility score of the text. The user interface is designed to be user-friendly, providing a clear and intuitive experience for users to interact with the system.

Please enter the model's response:

Known for its longstanding contribution to research in the UAE, MBZUAI is home to a vast team of experts across multiple disciplines. The institution is recognized for its comprehensive array of programs at both undergraduate and graduate levels primarily focusing on the theoretical aspects of artificial intelligence.

Submit

Time elapsed: 6 seconds

Figure 6: The user interface of LOKI, submission page.

The screenshot displays the LOKI result page for a fact-checking task. On the left, the original text is shown with numbered markers (1-8) corresponding to the claim. The central panel shows the claim: "MBZUAI has a vast team of experts." It lists 9 evidences, with 0 refutes, 8 contextualizations, and 1 support. A detailed view of the support evidence is shown below, including the source URL and reasoning. On the right, a summary table shows the overall factuality score and a breakdown of claim types.

Overall Factuality	
83.33333333333334%	
7	Detected claims
5	Well Supported
1	Conflict
0	Controversial

Figure 7: The user interface of LOKI, result page.