

 **VILLAIN at AVerImaTeC:**
Verifying Image–Text Claims via Multi-Agent Collaboration

Jaeyoon Jung^{♣◇} **Yejun Yoon**[♡] **Seunghyun Yoon**[♣] **Kunwoo Park**^{♣♡}

[♣]School of AI Convergence, Soongsil University

[◇]MAUM AI Inc.

[♡]Department of Intelligent Semiconductors, Soongsil University

[♣]Adobe Research, USA

{jaeyoonskr, yejun0382}@soongsil.ac.kr, syoon@adobe.com, kunwoo.park@ssu.ac.kr

Abstract

This paper describes VILLAIN, a multimodal fact-checking system that verifies image–text claims through prompt-based multi-agent collaboration. For the AVerImaTeC shared task, VILLAIN employs vision–language model agents across multiple stages of fact-checking. Textual and visual evidence is retrieved from the knowledge store enriched through additional web collection. To identify key information and address inconsistencies among evidence items, modality-specific and cross-modal agents generate analysis reports. In the subsequent stage, question–answer pairs are produced based on these reports. Finally, the Verdict Prediction agent produces the verification outcome based on the image-text claim and the generated question–answer pairs. Our system ranked first on the leaderboard across all evaluation metrics. The source code is publicly available at <https://github.com/ssu-humane/VILLAIN>.

1 Introduction

Motivated by the spread of misleading multimodal content and the need for automated verification, the AVerImaTeC shared task at the 9th FEVER Workshop focuses on verifying real-world image–text claims using external evidence (Cao et al., 2025), reflecting web-based fact-checking scenarios in the wild. Given a claim and its metadata, systems must retrieve relevant evidence from the web or from a provided document and image collection and assess the claim’s validity. Claims are labeled as Supported, Refuted, Not Enough Evidence, or Conflicting Evidence/Cherry-picking, and a prediction is considered correct only if both the label and the accompanying evidence are appropriate.

This paper describes  VILLAIN, developed by Team HUMANE, a fact-checking system designed to verify image–text claims using vision–language model (VLM) agents. The system employs agents operating across multiple stages of

automated fact-checking. The first two stages, Evidence Retrieval and Analysis, aim to identify evidence that facilitates veracity prediction. Using state-of-the-art embedding and rerankers, we retrieve textual and visual evidence from three knowledge stores, which are provided with the original dataset and further enriched through additional web content collection. To address noise and inconsistencies among evidence items, we employ three agents that generate analysis reports. The subsequent stage is Question–Answer Generation, in which an agent produces question–answer pairs by synthesizing the reports from the previous stage. Finally, in the Verdict Prediction stage, an agent predicts a verdict label together with a textual justification and the selected question–answer pairs. Our system achieved first place on the leaderboard, obtaining a veracity score of 0.546 on the test set and outperforming all other systems across the official evaluation metrics.

2 Related Works

AVeriTeC (Schlichtkrull et al., 2023) is a text-only fact-checking dataset that was used in the two previous shared tasks hosted by the FEVER workshop. At the Seventh FEVER Workshop (Schlichtkrull et al., 2024), Rothermel et al. (2024) proposed an LLM-based claim verification framework consisting of six stages built on GPT-4o, demonstrating the strong performance of closed-source models. In contrast, Yoon et al. (2024) introduced HerO based on open LLMs and highlighted the importance of evidence retrieval for verification accuracy. The subsequent task hosted at the Eighth FEVER Workshop imposed additional computational constraints (Akhtar et al., 2025). Under these constraints, CTU AIC proposed a two-stage pipeline based on long-context retrieval-augmented generation (Ullrich and Drchal, 2025a) to capture salient sentences and paragraphs. HerO 2 (Yoon et al.,

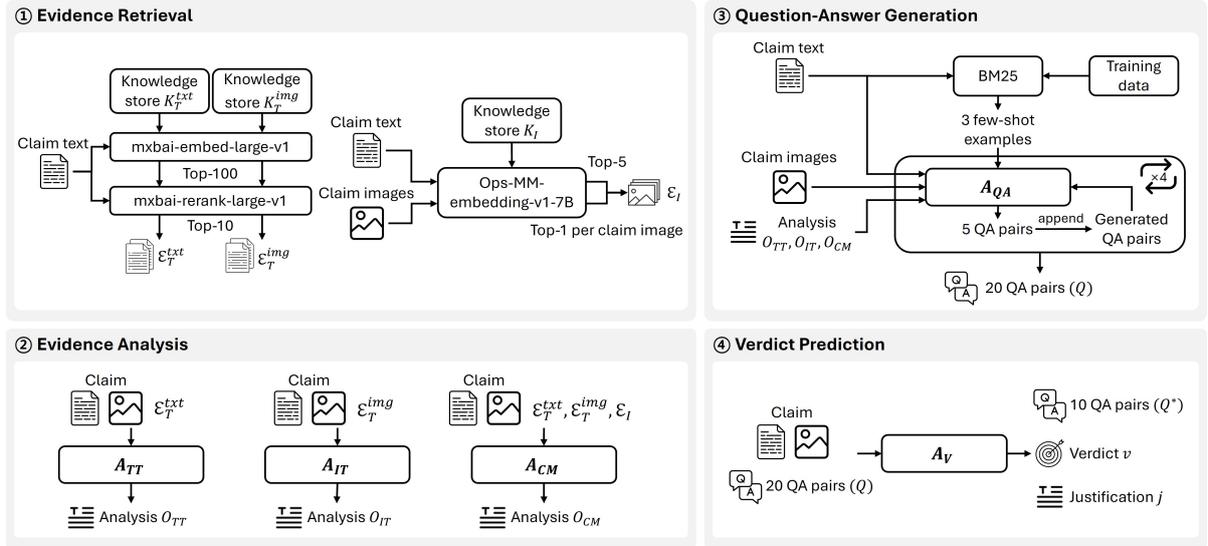


Figure 1: Overview of  VILLAIN.

2025) further improved upon HerO by applying model quantization and answer reformulation, emphasizing the importance of distilling answers to retain only the most relevant factual information.

VILLAIN is the first multi-agent system designed for AVerImaTeC, in which individual agents assume distinct roles across consecutive stages of the fact-checking process. Inspired by the strong systems proposed for earlier shared tasks, we adopt a prompt-based design that enables experimentation with both powerful open- and closed-source models, enhances the retrieval stage through reranking and knowledge store filling, and employs dedicated evidence analysis agents to identify and distill the most informative evidence for verification.

3 Task Description

The AVerImaTeC shared task focuses on developing automated multimodal fact-checking systems capable of verifying real-world image-text claims using retrieved web evidence. Given an image-text claim along with its associated metadata, a fact-checking system is required to generate a verdict. The supporting evidence may be retrieved either directly from the web or from the document and image evidence collections provided by the organizers. A response is considered correct only if both the predicted label is correct and the provided evidence is adequate. Specifically, a verdict is deemed correct only when the evidence score meets or exceeds the threshold $\lambda = 0.3$. Further details can be found in Cao et al. (2025).

4 Our System

This section describes  VILLAIN (Verifying Image-text cLaims via MuLti-Agent collaboratiON), developed by Team HUMANE, a multi-agent multimodal fact-checking pipeline that verifies claims through collaborative VLM agents. As illustrated in Figure 1, VILLAIN verifies a claim through a four-step sequential pipeline. Gemini-2.5-Pro (Comanici et al., 2025) is selected as the base model for all VLM agents.

4.1 Evidence Retrieval

The first stage aims to retrieve web evidence to verify a given claim. We use the set of knowledge stores provided with AVerImaTeC (Cao et al., 2025), denoted as $\mathcal{K} = \{\mathcal{K}_T^{txt}, \mathcal{K}_T^{img}, \mathcal{K}_I\}$. Each of the three stores contains a distinct type of evidence. Specifically, \mathcal{K}_T^{txt} denotes the knowledge store consisting of web text retrieved using a textual query; \mathcal{K}_T^{img} denotes the knowledge store comprising textual evidence obtained via reverse image search applied to the claim images; and \mathcal{K}_I denotes the knowledge store of web images retrieved using a textual query.

From a preliminary investigation of the data, we found that only URLs were provided for a substantial portion of samples in \mathcal{K}_T , while the corresponding text fields were missing. In addition, some evidence entries contained only generic website components (e.g., navigation bars, footers, or cookie notices) or access-restricted content (e.g., login walls) and therefore did not provide meaningful information for verification.

To address these issues, we perform a web content extraction process using Playwright¹, a browser automation framework, to populate knowledge-store entries that were previously empty or contained only generic or access-restricted content. A headless Chromium browser is employed to mitigate access restrictions that commonly affect non-browser clients. Knowledge-store entries are crawled either when their content is missing or when the existing content fails a usefulness check based on content length and the proportion of generic patterns. Table 1 reports statistics on the number of filled evidence items per claim before and after the crawling process. Further details can be found at Appendix B.

Textual Evidence Retrieval We retrieve textual evidence from both \mathcal{K}_T^{txt} and \mathcal{K}_T^{img} using mx-bai-embed-large-v1 (Lee et al., 2024) as the text embedding model and mx-bai-rerank-large-v1 (Shakir et al., 2024) as the text reranking model. For each knowledge store, we first retrieve the top-100 candidate evidence based on embedding similarity and then select top-10 evidence by reranking them. Following the winning approach of the previous AVeriTeC shared task (Ullrich and Drchal, 2025b), we segment each document into chunks of 2,048 characters. After retrieval, we augment each selected chunk with its adjacent chunks (one preceding and one following) to incorporate local context, yielding an evidence item consisting of three consecutive chunks. We denote the retrieved textual evidence sets as \mathcal{E}_T^{txt} and \mathcal{E}_T^{img} , respectively.

Visual Evidence Retrieval We retrieve visual evidence from \mathcal{K}_I using Ops-MM-embedding-v1-7B, a multimodal embedding model. Specifically, we retrieve the top-5 visual evidence items using the claim text as the query and retrieve the top-1 visual evidence item for each claim image. We denote the retrieved visual evidence set as \mathcal{E}_I .

4.2 Evidence Analysis

Some of the retrieved evidence may be irrelevant to the claim or lead to inconsistencies across different evidence items. To address these issues, we employ three VLM-based analysis agents, each responsible for analyzing a different set of retrieved evidence in the previous stage.

Specifically, the Text-Text agent (\mathcal{A}_{TT}) analyzes textual evidence associated with the claim text

¹<https://github.com/microsoft/playwright>

Split	Store	Status	Avg	Min	Max
Dev	\mathcal{K}_T^{txt}	Original	654.3	102	1,277
		Filled	837.5	133	1,683
	\mathcal{K}_T^{img}	Original	14.5	1	70
		Filled	16.3	1	91
Test	\mathcal{K}_T^{txt}	Original	1,429.0	160	2,900
		Filled	1,788.8	196	3,551
	\mathcal{K}_T^{img}	Original	23.5	1	174
		Filled	27.2	1	266

Table 1: Statistics of filled evidence per claim before and after URL content filling. Each URL with non-empty text content is counted as one evidence item.

(\mathcal{E}_T^{txt}), while the Image-Text agent (\mathcal{A}_{IT}) examines textual evidence derived from the claim images (\mathcal{E}_T^{img}). The Cross-Modal agent (\mathcal{A}_{CM}) investigates relationships across both textual and visual evidence (\mathcal{E}_I). The analysis process is formalized as follows:

$$\mathcal{O}_{TT} = \mathcal{A}_{TT}(c, \mathcal{E}_T^{txt}; \mathcal{P}_{TT}), \quad (1)$$

$$\mathcal{O}_{IT} = \mathcal{A}_{IT}(c, \mathcal{E}_T^{img}; \mathcal{P}_{IT}), \quad (2)$$

$$\mathcal{O}_{CM} = \mathcal{A}_{CM}(c, \mathcal{E}_T^{txt}, \mathcal{E}_T^{img}, \mathcal{E}_I; \mathcal{P}_{CM}), \quad (3)$$

where \mathcal{O} denotes the analysis output generated by each agent and \mathcal{P}_k ($k \in \{TT, IT, CM\}$) denotes the prompt used (shown in Figure 2, 3, and 4).

4.3 Question-Answer Generation

The Question-Answer (QA) generation agent \mathcal{A}_{QA} synthesizes the preliminary analysis outputs \mathcal{O}_{TT} , \mathcal{O}_{IT} , and \mathcal{O}_{CM} from the previous stage into QA pairs using a VLM. We generated five QA pairs iteratively by conditioning on the previously generated QA pairs. This design facilitates the generation of coherent follow-up question-answer pairs, while reducing redundancy by preserving context from the QA generation history.

Formally, the QA generation process at iteration i is defined as:

$$\mathcal{Q}^{(i)} = \mathcal{A}_{QA}(c, \mathcal{O}_{TT}, \mathcal{O}_{IT}, \mathcal{O}_{CM}, \mathcal{Q}^{(<i)}; \mathcal{P}_{QA}), \quad (4)$$

where \mathcal{P}_{QA} denotes the task-specific prompt for QA generation (shown in Figure 5), $\mathcal{Q}^{(i)}$ represents the QA pairs generated at iteration i , and $\mathcal{Q}^{(<i)}$ denotes the set of all QA pairs generated prior to iteration i . We employ in-context learning with retrieved three few-shot examples from the training data using BM25, and these examples are incorporated into the QA generation prompt \mathcal{P}_{QA} .

In our implementation, QA generation is conducted over four iterations, with five QA pairs generated at each iteration. Newly generated QA pairs are appended to the existing set and fed back into the model as part of the input for the next iteration, resulting in a total of 20 QA pairs. We denote the whole set of generated QA pairs as $\mathcal{Q} = \cup_{i=1}^4 \mathcal{Q}^{(i)}$.

4.4 Verdict Prediction

In the final stage, the Verdict Prediction agent \mathcal{A}_V verifies the claim. Given \mathcal{Q} , the set of QA pairs generated in the previous stage, \mathcal{A}_V is instructed to predict a verdict label along with a textual justification while selecting the top-10 most relevant QA pairs.

The verdict prediction process is formally defined as:

$$(v, j, \mathcal{Q}^*) = \mathcal{A}_V(c, \mathcal{Q}; \mathcal{P}_V), \quad (5)$$

where v denotes the predicted veracity label, j denotes the corresponding textual justification, \mathcal{P}_V denotes the prompt used (shown in Figure 6), and \mathcal{Q}^* denotes the subset of \mathcal{Q} selected by \mathcal{A}_V .

5 Evaluation Experiments

5.1 Experimental Setups

We evaluate model performance on the development set and conduct experiments with different vision–language model backbones, including the Gemini-2.5 (Comanici et al., 2025) and Qwen3-VL (Bai et al., 2025) model series, both of which incorporate inherent reasoning capabilities.

We evaluate our pipeline using Gemma-3-27B (Team et al., 2025), the model used for evaluation in the official AVerImaTeC shared task repository². Each evaluation is repeated five times to account for variability across runs, and we report the mean scores along with the standard deviation. As specified by the AVerImaTeC evaluation protocol, we report four metrics:

- **Q-Eval:** A reference-based recall metric for comparing the generated questions with the gold questions.
- **Evid-Eval:** A reference-based recall metric for comparing the system evidence with the gold evidence, computed via textual matching followed by image matching.

²https://github.com/abril4416/AVerImaTec_Shared_Task

System	Q-Eval	Evid-Eval	Veracity	Justification
HUMANE 🤖	0.923	0.583	0.645	0.543
ADA-AGGR	0.354	0.386	0.454	0.373
AIC CTU	0.822	0.347	0.375	0.304
<i>(Additional systems omitted for brevity)</i>				
Baseline	0.488	0.134	0.066	0.058

Table 2: Leaderboard results on the development set.

System	Q-Eval	Evid-Eval	Veracity	Justification
HUMANE 🤖	0.890	0.536	0.546	0.556
ADA-AGGR	0.370	0.463	0.537	0.433
AIC CTU	0.807	0.325	0.347	0.304
Baseline	0.555	0.171	0.114	0.132

Table 3: Leaderboard results on the test set.

- **Veracity:** The conditional accuracy of the predicted verdict label.
- **Justification:** A reference-based semantic similarity metric computed using Ev2R (Akhtar et al., 2024), which compares the generated justification with the human-annotated justification.

Both Veracity and Justification are reported conditionally, considering only instances for which the associated evidence score exceeds a predefined threshold λ . Instances falling below this threshold are treated as incorrect. We report Veracity and Justification score under both conditioned settings ($\lambda = 0.3$, the final evaluation setting of the shared task) and the unconditioned ($\lambda = 0$) settings.

We conducted our experiments on two machines: one equipped with eight H100 GPUs and 2 TB of RAM, and another with four NVIDIA A6000 GPUs and 256 GB of RAM. All experiments were performed in a computing environment configured with Python 3.11.13, PyTorch 2.7.0, Transformers 4.57.3, and SentenceTransformers 5.1.2.

5.2 Experimental Results

Development Set Results Table 2 presents the performance of VILLAIN on the development set, in comparison with the baseline and other competitive systems evaluated on the official Hugging Face leaderboard³. VILLAIN achieves the highest veracity accuracy of 0.645, followed by ADA-AGGR (0.454) and AIC CTU (0.375). The performance gap between VILLAIN and the baseline is substantial. Moreover, VILLAIN attains the best scores

³<https://huggingface.co/spaces/FEVER-IT/AVerImaTeC>

Model	Evidence Analysis	Q-Eval	Evid-Eval	Veracity
Gemini-2.5-Flash	✓	0.841 \pm .006	0.511 \pm .016	0.554 \pm .025
	✗	0.858 \pm .014	0.471 \pm .018	0.516 \pm .036
Qwen3-VL-32B	✓	0.803 \pm .009	0.489 \pm .017	0.532 \pm .022
	✗	0.848 \pm .011	0.422 \pm .017	0.483 \pm .023
Qwen3-VL-8B	✓	0.737 \pm .008	0.396 \pm .008	0.421 \pm .015
	✗	0.762 \pm .012	0.353 \pm .021	0.386 \pm .025

Table 4: Ablation study on the impact of evidence analysis agents. ✗ denotes disabling the evidence analysis agents and directly passing retrieved evidence to the QA generation agent, while ✓ denotes the full system with evidence analysis agents enabled.

across Q-EVAL, Evid-Eval, and Justification, indicating its effectiveness throughout multiple stages of the fact-checking pipeline.

Test Set Results Table 3 reports the test set performance of VILLAIN in comparison with the baseline and other top-ranked submissions. Our system achieves the highest veracity score (0.546), outperforming ADA-AGGR (0.537) and AIC CTU (0.347). Moreover, VILLAIN obtains the best scores across all evaluation metrics reflecting each stage of fact-checking.

Ablation Study: Evidence Analysis We conduct an ablation study to assess the contribution of the evidence analysis agents by removing them and directly passing the retrieved evidence to the QA generation agent. Table 4 compares model performance with and without evidence analysis agents. The results show that incorporating evidence analysis agents consistently improves the Evid-Eval scores across all evaluated models, with gains of 0.040 for Gemini-2.5-Flash, 0.067 for Qwen3-VL-32B, and 0.043 for Qwen3-VL-8B. These improvements suggest that the evidence analysis agents help identify and surface more relevant evidence, which in turn contributes to improved verdict prediction accuracy.

Ablation Study: Knowledge Store Filling We evaluate the impact of our filled knowledge store on task performance. Table 5 compares model performance when using the original knowledge store versus the filled knowledge store that has been.

Filling the knowledge store consistently improves Evid-Eval scores across all models, with gains of 0.03 for Gemini-2.5-Flash, 0.038 for Qwen3-VL-32B, and 0.036 for Qwen3-VL-8B. The improvements in question generation quality may be attributed to richer evidence, which can

Model	Fill	Q-Eval	Evid-Eval	Veracity
Gemini-2.5-Flash	✓	0.841 \pm .006	0.511 \pm .016	0.554 \pm .025
	✗	0.817 \pm .008	0.481 \pm .016	0.538 \pm .018
Qwen3-VL-32B	✓	0.803 \pm .009	0.489 \pm .017	0.532 \pm .022
	✗	0.782 \pm .005	0.451 \pm .010	0.505 \pm .015
Qwen3-VL-8B	✓	0.737 \pm .008	0.396 \pm .008	0.421 \pm .015
	✗	0.740 \pm .015	0.339 \pm .011	0.388 \pm .017

Table 5: Effect of knowledge store filling on model performance. ✓ denotes a filled knowledge store, and ✗ denotes the original knowledge store.

help models formulate more relevant diagnostic questions. Notably, Qwen3-VL-8B also shows a 0.033 improvement in verdict accuracy, indicating that additional evidence can help smaller models make more accurate predictions.

Ablation Study: Backbone VLM Comparison

Table 6 reports the experimental results on the development set obtained by varying the backbone VLMs used across all agent components, including evidence analysis, question generation, and verdict prediction. Specifically, we compare models from the Gemini-2.5 series (Pro and Flash) with those from the Qwen3-VL (32B, 8B, 4B, 2B) series.

Results demonstrate a clear performance trend in which Gemini-2.5-Pro outperforms all other models across most evaluation metrics. A single exception is observed for unconditioned verdict accuracy ($\lambda = 0$), where Qwen3-VL-32B performs comparably to Gemini-2.5-Pro and achieves a score that is 0.026 higher than Gemini-2.5-Flash. However, Qwen3-VL-32B yields lower scores for the other metrics, even when compared to Gemini-2.5-Flash, a more lightweight variant of Gemini-2.5-Pro.

6 Conclusion

We present  VILLAIN, the winning system in the AVerImaTeC shared task. It is a multimodal fact-checking pipeline that verifies image–text claims through prompt-based multi-agent collaboration. Using embedding-based retrieval and reranking, VILLAIN retrieves textual and visual evidence from the AVerImaTeC knowledge stores, which are further enriched with URL content filling. VLM agents perform modality-specific and cross-modal analysis, generate QA pairs, and produce a final verdict with justification and selected QA pair. Experimental results show that VILLAIN consistently outperforms competing systems across all evaluation metrics. Overall, these results highlight the effectiveness of multi-agent collaboration and itera-

Model	Size (B)	Q-Eval	Evid-Eval	Veracity		Justification	
				$\lambda = 0$	$\lambda = 0.3$	$\lambda = 0$	$\lambda = 0.3$
Gemini-2.5-Pro	-	0.913 \pm .009	0.588 \pm .014	0.908 \pm .000	0.634 \pm .023	0.757 \pm .011	0.566 \pm .014
Gemini-2.5-Flash	-	0.841 \pm .006	0.511 \pm .016	0.882 \pm .000	0.554 \pm .025	0.677 \pm .006	0.477 \pm .020
Qwen3-VL	32	0.803 \pm .009	0.489 \pm .017	0.908 \pm .000	0.532 \pm .022	0.650 \pm .010	0.441 \pm .020
	8	0.737 \pm .008	0.396 \pm .008	0.842 \pm .000	0.421 \pm .015	0.550 \pm .018	0.361 \pm .013
	4	0.713 \pm .003	0.382 \pm .018	0.862 \pm .000	0.420 \pm .032	0.540 \pm .007	0.321 \pm .020
	2	0.562 \pm .018	0.239 \pm .009	0.763 \pm .000	0.272 \pm .014	0.333 \pm .017	0.173 \pm .010

Table 6: Performance comparison between Qwen3-VL models of different scales and Gemini-2.5 models.

tive reasoning adopted in our system. By releasing our source code with support for multiple VLM backbones, we hope that VILLAIN will serve as a strong foundation for future research on reliable and transparent multimodal fact-checking systems.

Limitations

Generalizability While VILLAIN achieves state-of-the-art performance on AVerImaTeC, it has not been evaluated on other datasets. Therefore, the effectiveness of the system should be interpreted within the specific scope of the task and dataset. Future work could extend the dedicated approach into a more generalizable multi-agent framework for automated fact-checking.

Computational Costs For each claim, VILLAIN requires eight calls to a reasoning-capable VLM, which substantially increases the number of input and output tokens and, consequently, inference costs. Following recent advancements in efficient deployment (Feng et al., 2025) and the potential of small models for multi-agent systems (Belcak et al., 2025), future research could investigate a lightweight variant of the system.

Acknowledgments

This research was supported by the IITP (Institute of Information & Communications Technology Planning & Evaluation), funded by the Korea government (MSIT) (IITP-2026-RS-2022-00156360, IITP-2026-RS-2024-00430997, IITP-2026-RS-2020-II201602). Jaeyoon Jung and Yejun Yoon contributed equally to this work as co-first authors. Kunwoo Park is the corresponding author.

References

Mubashara Akhtar, Rami Aly, Yulong Chen, Zhenyun Deng, Michael Schlichtkrull, Chenxi Whitehouse,

and Andreas Vlachos. 2025. [The 2nd automated verification of textual claims \(AVeriTeC\) shared task: Open-weights, reproducible and efficient systems](#). In *Proceedings of the Eighth Fact Extraction and VERification Workshop (FEVER)*, pages 201–223, Vienna, Austria. Association for Computational Linguistics.

Mubashara Akhtar, Michael Schlichtkrull, and Andreas Vlachos. 2024. Ev2r: Evaluating evidence retrieval in automated fact-checking. *arXiv preprint arXiv:2411.05375*.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.

Peter Belcak, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. 2025. Small language models are the future of agentic ai. *arXiv preprint arXiv:2506.02153*.

Rui Cao, Zifeng Ding, Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2025. Averi-matec: A dataset for automatic verification of image-text claims with evidence from the web. *arXiv preprint arXiv:2505.17978*.

Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.

Sicheng Feng, Gongfan Fang, Xinyin Ma, and Xinchao Wang. 2025. Efficient reasoning models: A survey. *TMLR*.

Sean Lee, Aamir Shakir, Darius Koenig, and Julius Lipp. 2024. [Open source strikes bread - new fluffy embeddings model](#).

Mark Rothmel, Tobias Braun, Marcus Rohrbach, and Anna Rohrbach. 2024. Infact: A strong baseline for automated fact-checking. In *Proceedings of the*

Seventh Fact Extraction and VERification Workshop (FEVER), pages 108–112.

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 Fact Extraction and VERification Workshop (FEVER)*, pages 1–26, Miami, Florida, USA. Association for Computational Linguistics.

Michael Schlichtkrull, Zhijiang Guo, and Andreas Vlachos. 2023. [Averitec: A dataset for real-world claim verification with evidence from the web](#). *Advances in Neural Information Processing Systems*, 36:65128–65167.

Aamir Shakir, Darius Koenig, Julius Lipp, and Sean Lee. 2024. [Boost your search with the crispy mixedbread rerank models](#).

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, and 1 others. 2025. [Gemma 3 technical report](#). *arXiv preprint arXiv:2503.19786*.

Herbert Ullrich and Jan Drchal. 2025a. [Aic ctu@ fever 8: On-premise fact checking through long context rag](#). In *Proceedings of the Eighth Fact Extraction and VERification Workshop (FEVER)*, pages 274–280.

Herbert Ullrich and Jan Drchal. 2025b. [AIC CTU@FEVER 8: On-premise fact checking through long context RAG](#). In *Proceedings of the Eighth Fact Extraction and VERification Workshop (FEVER)*, pages 274–280, Vienna, Austria. Association for Computational Linguistics.

Yejun Yoon, Jaeyoon Jung, Seunghyun Yoon, and Kunwoo Park. 2024. [HerO at AVeriTeC: The herd of open large language models for verifying real-world claims](#). In *Proceedings of the Seventh Fact Extraction and VERification Workshop (FEVER)*, pages 130–136, Miami, Florida, USA. Association for Computational Linguistics.

Yejun Yoon, Jaeyoon Jung, Seunghyun Yoon, and Kunwoo Park. 2025. [Team HUMANE at AVeriTeC 2025: HerO 2 for efficient fact verification](#). In *Proceedings of the Eighth Fact Extraction and VERification Workshop (FEVER)*, pages 224–228, Vienna, Austria. Association for Computational Linguistics.

A Model Specifications

We evaluate models from the Qwen3-VL series (thinking variants) as well as the Gemini-2.5 series.

- **Qwen3-VL-2B**: <https://huggingface.co/Qwen/Qwen3-VL-2B-Thinking>
- **Qwen3-VL-4B**: <https://huggingface.co/Qwen/Qwen3-VL-4B-Thinking>
- **Qwen3-VL-8B**: <https://huggingface.co/Qwen/Qwen3-VL-8B-Thinking>
- **Qwen3-VL-32B**: <https://huggingface.co/Qwen/Qwen3-VL-32B-Thinking>
- **Gemini-2.5-Flash**: <https://ai.google.dev/gemini-api/docs/models?hl=en#gemini-2.5-flash>
- **Gemini-2.5-Pro**: <https://ai.google.dev/gemini-api/docs/models?hl=en#gemini-2.5-pro>

For the embedding and reranking components, we use the following models:

- **mxbai-embed-large-v1**: <https://huggingface.co/mixedbread-ai/mxbai-embed-large-v1>
- **mxbai-rerank-large-v1**: <https://huggingface.co/mixedbread-ai/mxbai-rerank-large-v1>
- **Ops-MM-embedding-v1-7B**: <https://huggingface.co/OpenSearch-AI/Ops-MM-embedding-v1-7B>

For the Gemma 3 27B model used in the evaluation, we employ:

- **Gemma-3-27b**: <https://huggingface.co/google/gemma-3-27b-it>

B Knowledge Store Construction

The AVerImaTeC shared task provides a knowledge store containing pre-retrieved evidence for each claim. However, we observed that many evidence entries had empty `url2text` fields, limiting the textual information available for fact-checking. To address this, we developed an automated URL content extraction pipeline to fill these missing fields.

B.1 URL Content Extraction Methodology

We implemented a robust web scraping pipeline that handles diverse web content types and JavaScript-heavy websites. The pipeline employs a two-tier approach:

Standard HTTP Requests For static websites, we use the `requests`⁴ library with custom headers to fetch HTML content. This approach is efficient for news sites, academic publications, and most standard web pages.

JavaScript Rendering For dynamic websites that require JavaScript execution (e.g., social media platforms and archive sites), we employ Playwright with Chromium to fully render pages before extracting content. Requests to domains such as Facebook, Instagram, Twitter/X, TikTok, LinkedIn, Quora, YouTube, and web archive services are automatically routed to Playwright, as these sites typically fail to return complete content when accessed via standard HTTP requests.

We filter out login walls, navigation elements, and generic footer content to ensure only substantive information is retained.

C Prompts Used

This section presents the complete prompts used by each agent in the 🦹 VILLAIN pipeline, together with concrete examples drawn from actual system outputs.

⁴<https://pypi.org/project/requests/>

Role

You are an expert AI Fact-Checker. Your specific task is to verify the **textual assertions** of a claim using the provided text-based sources.

Input Data

1. The Claim (Target for Verification)

- **Claimant (Speaker):** Unknown
- **Claim Date:** 2024-01-15
- **Claim Text:** Photo shows a 101-year-old woman who has given birth to her 17th child.
- **Claim Images:** [Image of elderly woman holding baby]

(Note: Use images for context, but focus verification on the text.)

2. Retrieved Evidence

- Retrieved Text Sources:

- [1] (snopes.com) On 7 January 2016, the web site World News Daily Report published an article reporting that a 101-year-old Italian woman gave birth... As with all of World News Daily Report's content, the claim was completely fabricated. The article swiped an unrelated image from a Reddit thread...
- [2] (thequint.com) The 101-year-old woman seen in the photo is Rosa Camfield, holding her great-granddaughter Kaylee, and not her own child as claimed. Camfield passed away in 2015...
- [3] (vishvasnews.com) Vishvas News investigated and found that the claim is false... The article headlined '101-Year-Old Great-Grandmother in Heartwarming Viral Photo Dies' carries the same image...

Instructions

1. **Contextual Understanding:** Analyze the Claim Text in conjunction with the Claim Images to fully understand the user's intent.
2. **Textual Verification:** Compare the factual claims made in the **Claim Text** against the **Retrieved Text Sources**. Look for: Factual support (dates, names, events); Contradictions or logical fallacies.
3. **Identify Information Gaps:** Explicitly state what information is missing from the *sources* that is needed to verify the text claim fully.

Output Format

1. Key Verification Facts

- * [Fact]: (Evidence from sources supporting/refuting the text claim)

2. Missing Information

- * [Gap]: (Crucial info missing from sources)

3. Analysis

(Summary of how text sources align with the claim text)

Model Output:

1. Key Verification Facts

- * [Fact]: The woman in the photograph is Rosa Camfield, who was 101 years old when the photo was taken. The baby is her great-granddaughter, Kaylee. (The Quint, Vishvas News)
- * [Fact]: The story of a 101-year-old Italian woman named Anatolia Vertadella giving birth originated from a website called "World News Daily Report," which has a disclaimer stating its articles are satirical and their content is fictional. (Snopes)
- * [Fact]: The official Guinness World Record for the oldest mother is Maria del Carmen Bousada Lara, who gave birth at age 66, making the claim of a 101-year-old giving birth medically implausible. (World Record Academy, CNN)
- * [Fact]: The woman in the photo, Rosa Camfield, passed away in March 2015, shortly after the photo went viral. (The Quint, Vishvas News)

2. Missing Information

- * The provided sources are sufficient to verify the claim. They identify the origin of the false story and the true context of the photograph.

3. Analysis

The claim that the photo shows a 101-year-old woman who has given birth to her 17th child is **false**. The image is authentic but miscaptioned—it shows 101-year-old Rosa Camfield holding her two-week-old great-granddaughter, Kaylee, in 2015. The narrative of "Anatolia Vertadella" is a piece of fiction from the satirical website "World News Daily Report."

Figure 2: Prompt used for \mathcal{A}_{TT} with an example of the input and output. Blue text indicates actual model output.

Role

You are an expert AI Fact-Checker. Your specific task is to verify the **visual content** of the claim using the provided text-based sources.

Input Data

1. The Claim (Target for Verification)

- **Claimant (Speaker):** Unknown
- **Claim Date:** 2024-01-15
- **Claim Text:** Photo shows a 101-year-old woman who has given birth to her 17th child.
- **Claim Images:** [Image of elderly woman holding baby]

(Note: Use claim text to understand what the image purports to show.)

2. Retrieved Evidence

- Retrieved Text Sources:

- [1] (africacheck.org) The photo shows Rosa Camfield, 101, cradling her newborn great-granddaughter. Camfield died just weeks later, on 30 March...
- [2] (chequeado.com) La foto corresponde a una mujer con su bisnieta y no es actual sino de 2015. La verdadera mujer de la foto es Rosa Camfield...
- [3] (snopes.com) World News Daily Report regularly publishes outlandish and false claims alongside repurposed photographs...

Instructions

1. **Visual Analysis:** Analyze the visual elements in the **Claim Images** (landmarks, people, signs, weather).
2. **Cross-Modal Verification:** Check if the events or descriptions in the **Retrieved Text Sources** explain or contradict the visual elements.
Example: Does the text report mention the specific objects or environment seen in the images?
3. **Identify Gaps:** What visual details are not explained by the text sources?

Output Format

1. Visual-Text Corroboration

- * [Point]: (How text sources confirm/deny specific visual elements)

2. Missing Context

- * [Gap]: (Visual details not mentioned in the text sources)

3. Analysis

(Summary of the consistency between the image and the text reports)

Model Output:

1. Visual-Text Corroboration

- * [Confirmed Visual Element - Elderly Woman]: The image shows a very elderly woman, consistent with the claim's description of a 101-year-old. Sources (The Quint, Snopes, Africa Check, Chequeado) confirm the woman is Rosa Camfield, who was 101 years old when the photo was taken in 2015.
- * [Confirmed Visual Element - Newborn Baby]: The image clearly shows a newborn baby. Sources confirm the baby is Kaylee, the woman's great-granddaughter, who was two weeks old at the time.
- * [Contradicted Relationship]: The claim states the woman is the baby's mother. Sources from Snopes, The Quint, Africa Check, and Chequeado all state this is false—Rosa Camfield is the baby's great-grandmother.
- * [Contradicted Narrative]: The story about "Anatolia Vertadella" giving birth is identified by Snopes and Chequeado as a complete fabrication from "World News Daily Report," a known satirical/fake news website.

2. Missing Context

- * [True Identity]: The claim omits that the woman is Rosa Camfield and the baby is her great-granddaughter, Kaylee.
- * [Original Context]: The photo was taken in March 2015 in the United States and was shared by the baby's mother, Sarah Hamm, capturing the moment the 101-year-old met her great-granddaughter for the first time.

3. Analysis

The visual elements superficially match the claim's description, but the sources reveal the image has been entirely decontextualized. The photograph shows a touching moment between a great-grandmother and her great-granddaughter, not a post-birth scene. The visual content does not support the narrative presented in the claim.

Figure 3: Prompt used for \mathcal{A}_{IT} with an example of the input and output. Blue text indicates actual model output.

Role

You are an expert AI Fact-Checker. Your specific task is to analyze the **cross-modal relationships** of the claim using the provided text and image sources.

Input Data

1. The Claim (Target for Verification)

- **Claimant (Speaker):** Unknown
- **Claim Date:** 2024-01-15
- **Claim Text:** Photo shows a 101-year-old woman who has given birth to her 17th child.
- **Claim Images:** [CLAIM_IMG_1]

2. All Retrieved Evidence

- Retrieved Text Sources:

- [1] (snopes.com) On 7 January 2016, the web site World News Daily Report published an article...
- [2] (thequint.com) The 101-year-old woman seen in the photo is Rosa Camfield...
- [3] (worldrecordacademy.org) Anatolia Vertadella, 101 year old, has become the oldest mother in history...

- Retrieved Source Images: [RETRIEVED_IMG_1], [RETRIEVED_IMG_2], ...

(Note: [CLAIM_IMG_n] tags represent claim images, and [RETRIEVED_IMG_n] tags represent retrieved source images. Treat these tags as placeholders for the actual visual data.)

Instructions

1. **Source-to-Source Text Analysis:** Compare the retrieved text sources. Do they agree on key facts (dates, locations, names)? Identify any contradictions between sources.
2. **Cross-Modal Alignment:** Analyze if the **Retrieved Source Images** align with the narratives in the **Retrieved Text Sources**.
Example: If Text Source A describes a “sunny protest,” does Image Source B show a sunny environment?
3. **Global Narrative Reconstruction:** Synthesize a coherent timeline or event description based on *all* available evidence.
4. **Reliability Assessment:** Identify if any source seems like an outlier or low-quality compared to others.

Output Format

1. Evidence Consistency Check

- * [Text-Text]: (Do text sources agree? Note contradictions.)
- * [Image-Text]: (Do source images support the source texts?)
- * [Image-Image]: (Is there visual consistency between the claim image and sources, and among sources themselves?)

2. Global Context Summary

(A unified summary of the event based on the combined evidence, independent of the user’s claim)

3. Conflict Alert

- * [Conflict]: (Critical discrepancies between sources, if any)

Model Output:

1. Evidence Consistency Check

- * [Text-Text]: There is a major contradiction. Group 1 (World Record Academy, guardian.ng, etc.) repeats the story of “Anatolia Vertadella” giving birth. Group 2 (Snopes, The Quint, Vishvas News, Africa Check) debunks this as a fabrication from “World News Daily Report,” a satirical website. The fact-checkers identify the woman as Rosa Camfield and the baby as her great-granddaughter.
- * [Image-Text]: Direct conflict. Group 1 mislabels the image as “Anatolia Vertadella” with her newborn son. Group 2 correctly identifies it as Rosa Camfield meeting her great-granddaughter.
- * [Image-Image]: The claim image is visually identical to retrieved images, confirming all sources discuss the same photograph.

2. Global Context Summary

In March 2015, a photograph was taken of 101-year-old Rosa Camfield holding her two-week-old great-granddaughter, Kaylee. The photo went viral. In January 2016, the satirical website “World News Daily Report” created a fabricated story about “Anatolia Vertadella” giving birth at 101, using this unrelated photo. Multiple fact-checking organizations have debunked this story.

3. Conflict Alert

- * [Conflict]: Critical conflict between the claim and verified facts. Identity Misrepresentation: The woman is Rosa Camfield (American), not “Anatolia Vertadella” (Italian). Relationship Misrepresentation: The baby is her great-granddaughter, not her child. Event Fabrication: The birth story is fiction from a satirical news source.

Figure 4: Prompt used for \mathcal{A}_{CM} with an example of the input and output. Blue text indicates actual model output.

Role

You are the Lead Fact-Checking Adjudicator. Your task is to synthesize preliminary analyses into **decisive Question-Answer (QA) pairs** that consolidate the key evidence and reasoning required to form a final verdict.

Input Data

1. The Claim (Target for Verification)

- **Claimant (Speaker):** Unknown
- **Claim Date:** 2024-01-15
- **Claim Text:** Photo shows a 101-year-old woman who has given birth to her 17th child.
- **Claim Images:** [Image attached]

2. Preliminary Analyses

[Agent 1 Output]: Key Verification Facts: The woman in the photograph is Rosa Camfield... The story originated from World News Daily Report...

[Agent 2 Output]: Visual-Text Corroboration: The image shows a very elderly woman with a newborn baby... Missing Context: True identity of the individuals...

[Agent 3 Output]: Evidence Consistency Check: Text-Text shows major contradiction... Global Context Summary: In March 2015, Rosa Camfield...

3. Few-shot Learning Examples

(Retrieved via BM25 from training set based on claim similarity)

4. Previously Generated QA Pairs

(Do NOT repeat these or ask similar questions)

[Previously generated pairs listed here in subsequent iterations]

Instructions

Synthesize Diagnostic QAs (The Reasoning Basis)

Analyze the provided forensic reports to extract the **core information** necessary to predict the verdict. Formulate **N high-impact QA pairs** that:

- **Isolate Key Evidence:** Focus on dates, locations, inconsistencies, or manipulation traces that act as “smoking guns.”
- **Resolve Ambiguity:** Ask and answer questions that clarify whether the evidence is sufficient or conflicting.
- **Serve as Proof:** Each QA must act as a logical premise supporting your final decision.

Output Format (JSON Only)

```
{"qa_pairs": [{"question": "<Question 1>", "answer": "<Full statement answer 1>"}, {"question": "<Question 2>", "answer": "<Full statement answer 2>"}]}
```

Model Output:

```
{"qa_pairs": [{"question": "Who are the people in the photograph and what is their actual relationship?", "answer": "The photograph shows 101-year-old Rosa Camfield holding her great-granddaughter, Kaylee, in March 2015. The baby is not her child."}, {"question": "Is it medically plausible for a 101-year-old to give birth?", "answer": "No, giving birth at 101 is considered medically implausible. The verified Guinness World Record for the oldest mother is Maria del Carmen Bousada Lara, who gave birth at the age of 66."}, {"question": "What is the origin of the story about a 101-year-old Italian woman named Anatolia Vertadella giving birth?", "answer": "The story is a complete fabrication that originated from 'World News Daily Report,' a website known for publishing satirical and fictional content."}, {"question": "When did the woman in the photo, Rosa Camfield, pass away?", "answer": "Rosa Camfield passed away in March 2015, shortly after the photo was taken, which was years before many versions of the false claim began circulating."}]}
```

Figure 5: Prompt used for \mathcal{A}_{QA} with an example of the input and output. Blue text indicates actual model output in JSON format.

Role

You are the Lead Fact-Checking Adjudicator. Your task is to select the most relevant QA pairs, assess veracity, and provide a final verdict with justification.

Input Data

1. The Claim (Target for Verification)

- **Claimant (Speaker):** Unknown
- **Claim Date:** 2024-01-15
- **Claim Text:** Photo shows a 101-year-old woman who has given birth to her 17th child.
- **Claim Images:** [Image attached]

2. Generated Question-Answer Pairs

1. **Q:** Who are the people in the photograph?
A: The photograph shows 101-year-old Rosa Camfield holding her great-granddaughter, Kaylee, in March 2015.
2. **Q:** What is the origin of the story about Anatolia Vertadella?
A: The story is a complete fabrication from "World News Daily Report," a satirical website.
3. **Q:** Is it medically plausible for a 101-year-old to give birth?
A: No. The Guinness World Record for oldest mother is 66 years old.
4. **Q:** What happened to Rosa Camfield after the photo was taken?
A: Rosa Camfield passed away on March 30, 2015, just weeks after meeting her great-granddaughter.

...

Instructions

1. **Select Best QA Pairs:** From the generated QA pairs above, select the **N most relevant and informative** pairs for verification.
2. **Determine Verdict:** Choose the single best label:
 - **Supported**
 - **Refuted**
 - **Not Enough Evidence**
 - **Conflicting Evidence/Cherry-picking**
3. **Write Justification:** A cohesive summary explaining the verdict based on the selected QA pairs.
4. **JSON Output:** Output **ONLY** a valid JSON object matching the format below.

Output Format (JSON Only)

```
{"questions": [{"question": "<Selected question 1>", "answer": "<Answer 1>"}, ...],  
"veracity_verdict": "<Label>", "justification": "<Explanation>"}
```

Model Output:

```
{  
  "questions": [  
    {"question": "Who are the people in the photograph and what is their actual relationship?",  
      "answer": "The photograph shows 101-year-old Rosa Camfield holding her great-granddaughter, Kaylee, in March 2015. The baby is not her child."},  
    {"question": "What is the origin of the story about a 101-year-old Italian woman named Anatolia Vertadella giving birth?",  
      "answer": "The story is a complete fabrication that originated from 'World News Daily Report,' a website known for publishing satirical and fictional content."},  
    {"question": "Is it medically plausible for a 101-year-old to give birth?",  
      "answer": "No, giving birth at 101 is considered medically implausible. The verified Guinness World Record for the oldest mother is Maria del Carmen Bousada Lara, who gave birth at the age of 66."}],  
  "veracity_verdict": "Refuted",  
  "justification": "The claim is a complete fabrication. The photograph does not show a mother and her newborn child, but rather 101-year-old Rosa Camfield holding her great-granddaughter, Kaylee, in March 2015. This was confirmed by her family, who originally shared the photo. Furthermore, the woman in the photo, Rosa Camfield, passed away in 2015, making it impossible for her to be the subject of a story that began circulating widely in 2016. The narrative of a 101-year-old Italian woman giving birth originated from 'World News Daily Report,' a well-known satirical website that explicitly states its content is fictional. The claim is also medically implausible, as the confirmed world record for the oldest mother is 66, not 101."  
}
```

Figure 6: Prompt used for \mathcal{A}_V with an example of the input and output. Blue text indicates the actual model output, including the final verdict of Refuted.