

# SciClaimEval: Cross-modal Claim Verification in Scientific Papers

Xanh Ho,<sup>1</sup> Yun-Ang Wu,<sup>2,3</sup> Sunisth Kumar,<sup>4</sup> Tian Cheng Xia,<sup>5\*</sup>  
Florian Boudin,<sup>6</sup> André Greiner-Petter,<sup>1,7</sup> and Akiko Aizawa,<sup>1,3,4</sup>

<sup>1</sup>National Institute of Informatics, Japan    <sup>2</sup>National Taiwan University    <sup>3</sup>NII LLMC, Japan

<sup>4</sup>The University of Tokyo, Japan    <sup>5</sup>University of Bologna, Italy

<sup>6</sup>Inria, LS2N, Nantes Université, France    <sup>7</sup>University of Göttingen, Germany

{xanh, yunangwu, aizawa}@nii.ac.jp, sunisth@g.ecc.u-tokyo.ac.jp

tiancheng.xia@studio.unibo.it, florian.boudin@univ-nantes.fr, greinerpetter@gipplab.org

## Abstract

We present SciClaimEval, a new scientific dataset for the claim verification task. Unlike existing resources, SciClaimEval features authentic claims, including refuted ones, directly extracted from published papers. To create refuted claims, we introduce a novel approach that modifies the supporting evidence (figures and tables), rather than altering the claims or relying on large language models (LLMs) to fabricate contradictions. The dataset provides cross-modal evidence with diverse representations: figures are available as images, while tables are provided in multiple formats, including images, LaTeX source, HTML, and JSON. SciClaimEval contains 1,664 annotated samples from 180 papers across three domains, machine learning, natural language processing, and medicine, validated through expert annotation. We benchmark 11 multimodal foundation models, both open-source and proprietary, across the dataset. Results show that figure-based verification remains particularly challenging for all models, as a substantial performance gap remains between the best system and human baseline.

**Keywords:** Claim Verification, Cross-modal, Scientific Papers

## 1. Introduction

Scientific claim verification involves determining whether claims made in research papers are supported or refuted by the accompanying evidence (Wadden et al., 2020; Guo et al., 2022). With the rapid rise of generative AI and large language models (LLMs), the number of submissions to scientific conferences and journals has grown substantially, creating a greater demand for tools that help reviewers assess the validity of authors' claims. Accurate and reliable claim verification systems could significantly strengthen the peer-review process by automatically identifying unsupported or inconsistent claims. However, the ability of existing models, including recent multimodal LLMs (MLLMs), to perform this task remains largely unexplored due to the lack of a comprehensive dataset.

Over the past few years, several datasets have been introduced for claim verification in scientific papers, including SciFact (Wadden et al., 2020), PubHealth (Kotonya and Toni, 2020), SciTab (Lu et al., 2023), and more recently, SciVer (Wang et al., 2025a). SciFact and PubHealth are text-only datasets in the medical domain, while SciTab focuses on computer science and provides tables as evidence. SciVer extends this line of work by incorporating multimodal evidence, while remaining focused on computer science. However, SciVer has two key limitations: first, its claims are synthetic, written by experts rather than drawn from actual papers, and second, its tables are provided solely as

images, without corresponding structured formats. More broadly, a persistent limitation across most existing datasets lies in how unsupported claims are constructed. These are often created by editing supported claims to introduce contradictions with the evidence, which can lead to artifacts and heuristic cues. For example, many refuted claims simply rely on inserting negation words such as not, making the task less realistic.

In this paper, we present SciClaimEval, a new dataset for cross-modal scientific claim verification. Our work differs from existing resources in three important ways. First, all claims, both supported and refuted, are authentic, sourced directly from published papers across three domains: machine learning (ML), natural language processing (NLP), and medicine. Second, unlike prior work that generate refuted claims by altering original statements or relying on LLMs, we introduce a novel strategy that creates negative examples by disturbing the supporting evidence itself, specifically by modifying figures and tables. Third, SciClaimEval provides cross-modal evidence with rich and diverse representations: both tables and figures are included, and tables are available in multiple formats, including images, LaTeX source, HTML, and JSON.

Specifically, SciClaimEval is constructed in three main steps. First, we collect a pool of papers spanning three domains: NLP, ML, and medicine. For the medical domain, we use papers from PeerJ, while for NLP and ML, we gather papers from arXiv, referencing accepted papers from the ACL Anthology and NeurIPS as guidance. For each paper, we

\*Research conducted during internship at NII, Japan.

Year	Dataset	Size	Domain	Fact Source (Positive)	Fact Source (Negative)	Authentic Claims	Input Context	Sources	Verdict
2021	SEM-TAB-FACTs	5,715	Multi	Crowd Original sentences	Crowd	XX	A table (.xml)	ScienceDirect	3 labels
2023	SciTab	1,224	CS	GPT followed by verification	Claim modification	✓X	A table (.json)	SciGen	3 labels
2025	MMSci-Eval	3,114	CS	GPT followed by verification	Claim modification	XX	A table image	SciGen	3 labels
2025	SciAtomicBench	2,568	Multi	GPT followed by verification	Claim modification	XX	A table (.json)	SciGen, PubTables-1M, Financials, MatSciTable	2 labels
2025	SciVer	3,000	CS	Experts	Claim modification	XX	Textual paras, multiple charts, multiple tables (table image)	arXiv papers	2 labels
2025	MuSciClaims	1,515	Multi	Original sentences	Claim modification	✓X	A figure	Nature Physics, Journal of the American Chemical Society, and Cell	3 labels
2025	SciClaimEval (ours)	1,664	Multi	Original sentences	Evidence modification	✓✓	A table* or a figure	arXiv, PeerJ	2 labels

Table 1: Comparison of SciClaimEval with existing multimodal scientific claim verification datasets. For the Authentic Claims column, the two bars from left to right represent a supported claim and a refuted claim, respectively. The asterisk (\*) next to Table indicates that our table is available in multiple formats, including table image, JSON, and HTML.

extract the main text, figures, and tables. Second, we perform claim–evidence pair extraction. Using keywords such as “Table 1” or “Tab. 1”, we identify sentences that explicitly reference a table or figure and pair each of them with the corresponding evidence. Third, we conduct expert annotation. All claim-evidence pairs are reviewed by experts, who carry out two tasks: 1) claim-evidence verification and 2) evidence modification. When a claim is annotated as supported, the annotators modify the corresponding evidence to create an unsupported version. In total, SciClaimEval contains 1,664 claims from 180 papers, covering both figures and tables across three domains.

We then evaluate 11 multimodal foundation models on SciClaimEval, covering both open-source and proprietary models of varying sizes. Our results show that the figure-based subset is challenging for all models, including o4-mini, with a substantial gap remaining between the best model performance and the human baseline. In contrast, the table-based subset is primarily useful for assessing open-source MLLMs, as o4-mini performs close to the human baseline. Additionally, our table data includes diverse formats, providing a valuable resource for future research on processing scientific papers.<sup>1</sup>

## 2. Related Work

In this section, we first review studies on general-domain claim verification, focusing on text-based

approaches, and then discuss research related to multimodal scientific claim verification.

**Claim Verification.** Claim verification (fact-checking) has long been a central problem in NLP and AI, with extensive progress surveyed by Guo et al. (2022). The task aims to determine whether a claim is supported, refuted, or unverifiable given available evidence. Research on general-purpose, text-based verification has primarily focused on two sources: news articles and Wikipedia. In the news domain, several benchmark datasets have been developed, including LIAR (Wang, 2017) and MultiFC (Augenstein et al., 2019). In the Wikipedia domain, widely used datasets include FEVER (Thorne et al., 2018) and HoVer (Jiang et al., 2020). Both datasets contain claims that often require reasoning over multiple documents for verification.

Beyond purely textual settings, the community has also introduced resources for verification over structured and multimodal evidence. These include tables (TabFact (Chen et al., 2020) and FEVEROUS (Aly et al., 2021)), figures and charts (ChartCheck (Akhtar et al., 2024)), and knowledge graphs (FactKG (Kim et al., 2023)), broadening the scope of reasoning beyond unstructured text.

In parallel, domain-specific efforts have targeted scientific claims, where veracity judgments depend on rigorously sourced scholarly evidence. Notable datasets include SciFact (Wadden et al., 2020), PubHealth (Kotonya and Toni, 2020), and SciFact-Open (Wadden et al., 2022).

<sup>1</sup>Our dataset is available at <https://sciclaimeval.github.io/>

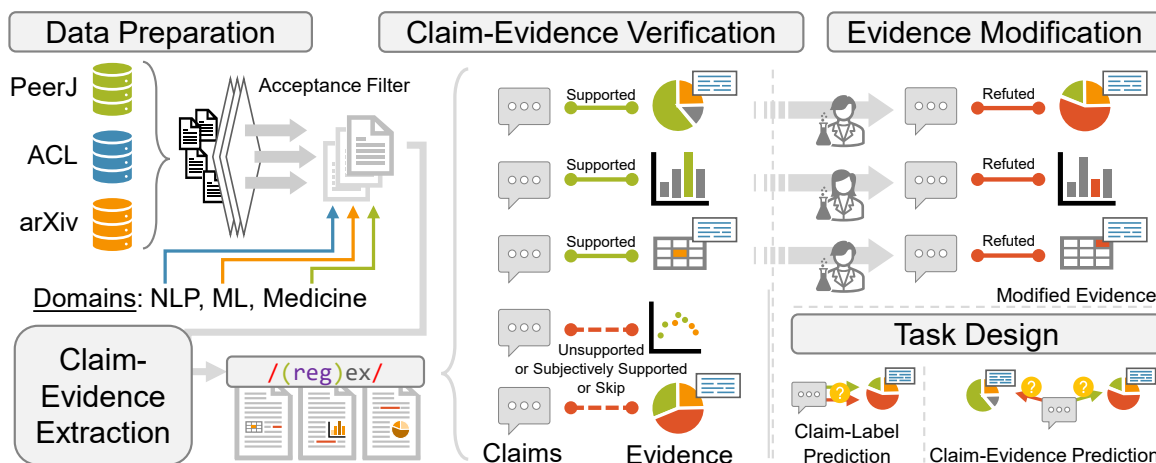


Figure 1: Our dataset construction pipeline consists of three main steps: data preparation, automatic claim-evidence extraction, and human annotation (Subsections 3.1, 3.2, and 3.3). The human annotation process involves two tasks: claim-evidence verification and evidence modification. After collecting all samples, we design two subtasks in our dataset: claim-label prediction and claim-evidence prediction. Details are in 3.7.

**Multimodal Scientific Claim Verification.** Unlike the datasets presented in the previous section, this section discusses datasets that involve multimodal scientific content. Table 1 provides a comparison of existing datasets and our dataset. SciTab (Lu et al., 2023), SEM-TAB-FACTs (Wang et al., 2021), and SciAtomicBench (Zhang et al., 2025) are datasets that focus on using tables from scientific papers. The tables in SciTab and SciAtomicBench are represented in JSON format, whereas the tables in SEM-TAB-FACTs are represented in XML format. MMSci-Eval (Yang et al., 2025b) also uses tables as evidence, but the tables are represented as images. MuSciClaims (Lal et al., 2025) focuses on claim verification using figures; their figures are often complex and contain multiple subcharts and tables. SciVer (Wang et al., 2025a) emphasizes the use of multiple charts, tables, or textual paragraphs as evidence, highlighting reasoning over multiple pieces of information. A limitation of all existing datasets is that they only modify claims to create refuted examples. These modifications typically involve generating the opposite meaning or performing semantic flips of existing claims, which can lead to spurious patterns and shortcut reasoning.

Another key aspect of our dataset is that both supported and refuted claims are authentic claims extracted directly from scientific papers. While datasets such as SciTab and MuSciClaims also reuse authentic claims from existing papers, they do so only for supported claims, limiting the diversity of refuted examples. With AI increasingly assisting in scientific writing and reviewing, using authentic claims is crucial to reflect real-world challenges in scientific reasoning.

### 3. Dataset Construction and Analysis

In this section, we first describe our dataset construction process, which consists of three main steps: data preparation, automatic claim and evidence extraction, and human annotation, corresponding to Subsections 3.1, 3.2, and 3.3. We then provide detailed information about the resulting dataset, including dataset statistics, dataset analysis, dataset validation, and task design, presented in Subsections 3.4, 3.5, 3.6, and 3.7. Figure 1 illustrates the overall process of our dataset construction.

#### 3.1. Data Preparation

In our dataset creation process, we collect papers from three sources: PeerJ for the medical domain, ACL for NLP, and AI/ML conferences for machine learning. For NLP domain, papers are collected from the ACL Anthology and manually mapped to their arXiv versions when available. We select papers from a variety of conferences, including but not limited to EMNLP, ACL, EACL, NAACL, and IJCNLP. For AI/ML domain, we select arXiv papers that include comments indicating acceptance at AI/ML conferences, such as NeurIPS. The arXiv IDs from both the NLP and AI/ML collections are used to retrieve data from ar5iv (Ginev, 2024), which provides HTML-rendered versions of arXiv papers. The main text of each paper is cleaned and parsed into JSON format, storing the title, abstract, and a list of paragraphs for each section. Figures are extracted from ar5iv and saved in PNG format. To obtain high-quality table data, we additionally download the LaTeX source of each arXiv paper and extract tables directly from the source files, pre-

serving them in LaTeX format. For medical domain, we select papers published in PeerJ between 2024 and 2025 in the Medicine Articles category. Since PeerJ provides HTML versions of papers, we crawl the raw HTML and parse it into JSON. Figures and tables are extracted separately, with tables stored in HTML format and figures in PNG.

### 3.2. Claim-Evidence Extraction

We use regular expressions to identify paragraphs in each paper that mention either a table or a figure. When such a reference is found, we extract the corresponding paragraph, split it into individual sentences, and pair each sentence with the mentioned table or figure, treating it as supporting evidence. By applying this process to all papers, we construct a set of sentence–evidence pairs, where the evidence corresponds to either a table or a figure. Our dataset focuses solely on the main text of the papers, excluding any content from appendices, including their tables and figures. Following the findings of Ho et al. (2025), which suggest that a lack of context can lead to task ambiguity, we also include the preceding sentences from the same paragraph as a short contextual field for each claim sentence.

### 3.3. Human Annotation

Our human annotation process includes two tasks: claim-evidence verification and evidence modification. We first describe the annotators, then detail the annotation process.

**Annotators Information.** We recruited annotators from the students of our laboratory. Including the authors of this paper, we have a total of 11 annotators, all of whom are either graduate students or expert researchers in the NLP and ML domains. We did not provide direct monetary compensation; however, the annotators were rewarded with a fully funded travel trip covering all associated costs.

**Claim-Evidence Verification.** From the list of extracted claim–evidence pairs for each paper, annotators are presented with a claim, its corresponding evidence (either a figure or a table), and the preceding sentences from the same paragraph as short context. If the claim is the first sentence, the context is empty. Annotators are asked to: **1. Assign a label:** Choose one of the following: **Supported** (evidence clearly supports the claim with no ambiguity), **Subjectively Supported** (the claim contains subjective terms such as “large margin” or “competitive,” making it difficult for annotators to determine whether these adjectives are accurate), **Unsupported** (evidence does not support the claim), or

**Skip** (insufficient knowledge to judge, or the claim is problematic or purely descriptive). **2. Indicate context use:** Choose one of the following: No (the claim is understandable without context), Yes (the short context is needed), or Other sources (the full paper is needed to understand the claim).

**Evidence Modification.** If a claim–evidence pair is labeled as *Supported*, annotators proceed to the evidence modification task. We use the same input as in the previous task and provide annotators with a list of modification operations, including explanations and examples. The goal is to modify the table or figure so that the claim becomes *Unsupported* when paired with the altered evidence. For tables, annotators can perform the following operations:

- **Change the cell values:** Modify the content of one or more cells.
- **Swap rows or columns:** Move the name of a row or column to another one.
- **Alter the table:** Add or remove rows or columns.
- **Others:** Annotators may introduce additional changes.

For figures, annotators can perform the following operations:

- **Graph Flip:** Flip a graph or part of it.
- **Legend Swap:** Swap text in the legend.
- **Graph Swap:** Exchange graphs or subgraphs from the same paper.
- **Category Swap:** Swap category labels to contradict the claim.
- **Others:** Propose other types of modifications as appropriate for the sample.

In both cases, when annotators select **Others**, they are required to record the details of the changes.

### 3.4. Dataset Statistics

The statistics of our dataset are presented in Table 2. We report the number of claims in the validation and test sets, as well as for the entire dataset. Information is provided across different properties, including modality type (table or figure), context usage, and domain. We split the data into validation and test sets based on context use and the type of operation from the evidence modification task. All *Others* operations are included in the test set. Additionally, for tables, we retain the *Alter the table* operations and some *Swap rows or columns* operations in the test set.

Property	Val	Test	All Data
<i>Labels</i>			
#Supported	395	481	876
#Refuted	352	436	788
<i>Modality Type</i>			
Table	482	523	1,005
Figure	265	394	659
<i>Context Use</i>			
No	494	619	1,113
Short Context	148	190	338
Full Paper	105	108	213
<i>Domain</i>			
NLP	388	389	777
ML	162	240	402
Medicine	197	288	485
#Papers	139	164	180
Total examples	747	917	1,664

Table 2: Dataset statistics of SciClaimEval.

**Supported Claim Only.** As shown in Table 2, we have 876 supported claims but only 788 refuted claims. In theory, the numbers of supported and refuted claims should be the same, since each evidence table or figure is modified to create an unsupported claim using the altered evidence. However, during the dataset annotation process, we found that in some cases it was very difficult to modify the evidence in a way that remained logical and still posed a meaningful challenge to the models, rather than simply making superficial changes. As a result, we have 88 “supported claim only” samples, 67 of which come from figure-based evidence and 21 from table-based evidence.

**Different Formats of Table Evidence.** For figure evidence, the figures are provided in .png format. For table evidence, the original tables in the NLP and ML domains are in LaTeX format, while the tables in the Medicine domain are in .html format. We also obtain both .png and .json versions of the table data. To generate the .png format, the LaTeX files are first compiled into PDFs, which are then converted to PNG images using the Python pdf2image library (Belval, 2023). This library relies on built-in system tools for PDF manipulation. The HTML files are converted directly into images using wkhtmltoimage (Kulkarni, 2023), an open-source tool that renders HTML files into various formats. To estimate the accuracy of our table evidence in .png format, we randomly select 100 samples and manually evaluate them. We find that only one sample is rendered incorrectly because of overlapping columns. Additionally, three samples have issues with the caption, as two have incomplete captions

and one has no caption.

To generate JSON from LaTeX files, we first tried a rule-based approach but found it inadequate due to numerous edge cases. We then adopted GPT-5-nano to convert LaTeX and HTML tables into a predefined JSON format following the schema of Lu et al. (2023), which includes the table ID, caption, column names, and cell values. To estimate the accuracy of our table evidence in .json format, we randomly select 100 samples and manually evaluate them. We found 22 cases with minor issues (e.g., missing the top row due to multicolumns) and 19 cases with major issues where the JSON table content did not match the original table.

Information	Max	Min	Avg.
Claim length	91	7	25.6
Table caption length	134	3	29.9
Figure caption length	206	7	53.4
Context length	296	8	67.7

Table 3: Detailed analyses of text lengths (based on word count) in SciClaimEval.




Figure 2: Analyses of evidence-modifying operations in SciClaimEval.

### 3.5. Dataset Analysis

**Text Lengths.** Table 3 presents detailed analyses of the word counts for claims, table captions, figure captions, and short contexts. As shown in the table, figure captions are generally longer than table captions. On average, claims contain 25.6 words, ranging from 7 to 91 words.

**Evidence-Modifying Operations.** We present the operation analyses for table evidence and figure evidence modifications in Figure 2. As illustrated in the figure, the most common modification for table evidence is *changing cell values*, while for figure evidence, *legend swapping* is the most frequent. Annotators appear to be more creative when working with figure evidence, as they often select *others* as the operation type. In contrast, *others* is rarely chosen for table evidence. Upon

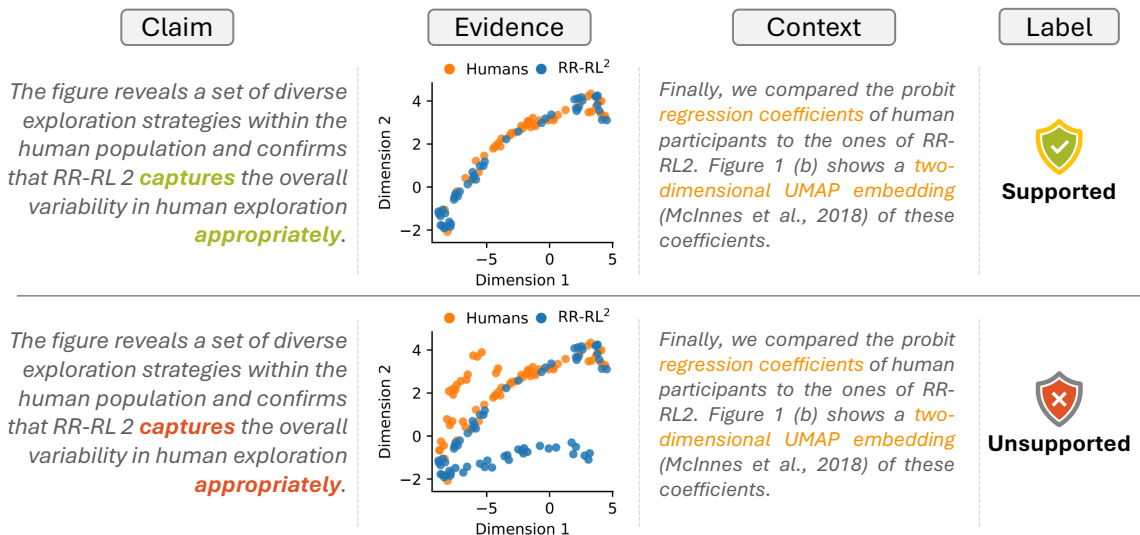


Figure 3: An example of an *others* modification in the figure evidence from our dataset involves creating an unsupported claim by adding spurious data points. The annotator labels this operation as *others*, with the specific detail noted as “adding fake data points.” The context provides the necessary information to understand the plot in the evidence.

examining the details categorized as *others* in the figure evidence, we find that annotators frequently perform operations such as *changing bar heights*, *manipulating data points* (e.g., moving or adding fake data points), *adjusting axes*, or *rearranging graphs*. Figure 3 shows an example in which the annotator labels the operation as *others*, with the specific detail noted as “adding fake data points.”

*swap* result in the most significant changes, which is reasonable since these two operations modify large areas of the chart. In contrast, *category swap* and *legend swap* produce minimal changes, as altering categories or legends is typically a localized operation. The *others* category lies between these two groups, as it contains a variety of mixed operations.




Figure 4: Violin plots showing the distribution of the Structural Similarity Index (SSIM; Wang et al., 2004) across five operation types. In each plot, the black dashed line indicates the group mean, and the numeric label to the right of the line denotes the corresponding average value.

#### Pixel Changes in Figure Evidence Modification.

Figure 4 illustrates the image similarity between the original and modified images for each operation type. The Structural Similarity Index (SSIM; Wang et al., 2004) was used as the metric for image similarity. We exclude 11 pairs with scaled or misaligned edits, as such cases would heavily skew the SSIM score. We observe that *graph flip* and *graph*

### 3.6. Dataset Verification

To establish a baseline for human performance on the dataset, we evaluated 80 samples, including 45 with table evidence and 35 with figure evidence. Each subset was independently annotated by two annotators, resulting in a total of four annotators. The annotators were graduate students and AI researchers. On the table subset, the average macro-F1 score was 87.9, with an inter-annotator agreement of 86.7%. For the figure subset, the average macro-F1 score was 89.6, with an agreement rate of 91.4%.

### 3.7. Task Design

After collecting all samples, we design two subtasks in our dataset. The first subtask is claim-label prediction, which is the main task and follows prior work on claim verification datasets. The input consists of a claim, an evidence file (which can be a figure or a table) along with its caption, and additional contextual information intended to reduce ambiguity in the sample. The output is a label indicating whether the claim is supported or refuted. The second subtask is claim-evidence prediction, whose goal is to identify which piece of evidence supports a given claim.

This task is particularly challenging because the two evidence files are highly similar, making it difficult to distinguish the correct supporting evidence.

## 4. Experiment

### 4.1. Experimental Settings

**Models.** For open-source multimodal LLMs, we use four variants of InternVL3\_5 (1B, 8B, 14B, and 38B) (Wang et al., 2025b); three variants of Qwen3-VL (4B, 8B, and 30B-A3B) (Yang et al., 2025a); two variants of LLaVA-v1.6 (llava-v1.6-mistral-7b and llava-v1.6-vicuna-13b) (Li et al., 2024); and Llama-3.2-11B-Vision (Grattafiori et al., 2024). We note that the instruction-tuned versions of these models are used. For the proprietary model, we use OpenAI o4-mini (OpenAI, 2025).

**Prompting Strategies.** Following the SciVer dataset (Wang et al., 2025a), we also employ zero-shot Chain-of-Thought (CoT; Wei et al., 2022) prompting in our experiments. As shown in Section 3 and Table 2, our dataset includes context usage information. **No** indicates that only the claim and a figure or table are needed. **Yes** indicates that a short context, consisting of the sentences preceding the claim, is required. **Full paper** indicates that the full text may be needed. Based on this, we design two setups: *no-context* for samples not requiring context and *use-context* for samples requiring either a short context or the full paper.

**Evaluation.** Following previous work on the scientific claim verification task (Lu et al., 2023; Ho et al., 2025), we use the macro-F1 evaluation metric in our experiments. However, since the task is a binary classification problem (Supported vs. Refuted), macro-F1 alone may be insufficient. A model may achieve a reasonable macro-F1 score through lucky guesses or by exploiting reasoning shortcuts or dataset biases. To mitigate this effect, we introduce a new evaluation metric, Pair Accuracy, defined as the number of correctly predicted pairs divided by the total number of pairs. A pair is considered correct only if both samples associated with the same claim, one Supported and one Refuted, are predicted correctly. This metric is more stringent: a robust model unaffected by biases should be able to correctly predict the labels for both evidence files corresponding to the same claim. Notably, while the random baseline for macro-F1 is 0.5, the random baseline for Pair Accuracy is 0.25, making Pair Accuracy a stricter and more discriminative metric than macro-F1.

### 4.2. Results

Table 4 presents the macro-F1 and pair accuracy of the models on our dataset.

**Macro-F1 vs. Pair Accuracy.** As shown in the table, Pair Accuracy scores are consistently lower than macro-F1 scores across all cases. This indicates that Pair Accuracy is a stricter evaluation metric, which helps reduce inflated performance caused by model guessing. Specifically, if a model correctly predicts the label for either the Supported or Refuted sample but fails on the other sample within the same claim pair, despite only slight changes in the evidence file, this suggests that the model does not truly understand the evidence. Instead, its predictions may rely on superficial or spurious features rather than genuine reasoning over the evidence.

**Using No Context vs. Context.** Based on the prompt strategies in Section 4.1, we evaluate each validation and test subset using two setups: no-context and use-context. As shown in the table, in most cases, models perform better without additional context than with context, with the no-context setting often achieving higher scores than settings that require context. Overall, these results suggest that solving the task without context is generally easier, although additional context can be beneficial in some cases. Since both tables and figures already include contextual information in their captions, providing extra context is not always helpful. When the task requires the model to jointly reason over both the context and the evidence, performance tends to decrease, indicating increased difficulty. In contrast, when context serves as supplementary information that does not require complex reasoning, it can improve performance.

**Open-source vs. Proprietary Models.** As shown in Table 4, o4-mini outperforms all open-source MLLMs across all settings, highlighting the performance gap that remains between proprietary and open-source models. Among open-source MLLMs, Qwen3-VL-30B-A3B achieves the best results in most cases, while Qwen3-VL-8B leads in the use-context test setting. As recent models, the Qwen3-VL series demonstrate clear progress in open-source MLLM development.

### 4.3. Analyses

To explore model performance in depth, we analyze predictions from several top-performing models. First, we compare samples using tables versus figures as evidence. Second, we examine the effect of evidence modifications on model behavior.

Model	Validation						Test					
	No Context		Use Context		Average		No Context		Use Context		Average	
	F1	P-Acc	F1	P-Acc	F1	P-Acc	F1	P-Acc	F1	P-Acc	F1	P-Acc
#Samples	494	231	253	121	747	352	619	294	298	142	917	436
llava-mistral-7b	45.9	1.3	49.8	2.5	47.9	1.7	47.4	2.4	49.7	0.7	48.4	1.8
llava-vicuna-13b	25.8	0.0	27.2	0.8	26.3	0.3	25.9	0.3	29.5	0.0	27.1	0.2
Llama-3.2-11B-Vision	49.0	11.7	47.8	9.1	48.6	10.8	50.1	17.3	44.8	8.5	48.5	14.4
InternVL3_5-1B	52.3	21.6	52.3	15.7	52.3	19.6	51.6	22.4	52.1	16.9	51.8	20.6
InternVL3_5-8B	68.5	40.7	59.0	24.8	65.4	35.2	66.4	38.8	55.2	26.1	63.0	34.6
InternVL3_5-14B	70.8	46.3	63.7	32.2	68.5	41.5	68.0	40.5	59.3	27.5	65.3	36.2
InternVL3_5-38B	70.8	45.0	61.8	30.6	67.8	40.1	70.1	45.6	64.0	34.5	68.2	42.0
Qwen3-VL-4B	71.5	46.3	68.7	38.8	70.6	43.8	70.4	45.6	67.5	36.6	69.6	42.7
Qwen3-VL-8B	72.2	47.6	71.3	45.5	72.1	46.9	70.5	46.6	<u>69.1</u>	<u>39.4</u>	70.1	44.3
Qwen3-VL-30B-A3B	<u>76.2</u>	<u>55.0</u>	<u>75.5</u>	<u>54.5</u>	76.0	<u>54.8</u>	<u>73.5</u>	<u>49.7</u>	67.0	<u>39.4</u>	<u>71.4</u>	<u>46.3</u>
o4-mini	<b>82.8</b>	<b>68.0</b>	<b>83.1</b>	<b>68.6</b>	<b>82.9</b>	<b>68.2</b>	<b>80.3</b>	<b>63.3</b>	<b>76.5</b>	<b>54.9</b>	<b>79.1</b>	<b>60.6</b>

Table 4: Macro-F1 (denoted as  $F1$  in the table) and pair accuracy (denoted as  $P\text{-Acc}$  in the table) of the models on our dataset. For the  $F1$  columns, the number in the *Samples* row represents the number of individual samples, whereas for the  $P\text{-Acc}$  columns, the number represents the number of sample pairs. Bold numbers indicate the best scores among different models, while underlined numbers represent the second-best scores.

Model	Validation		Test	
	F1	P-Acc	F1	P-Acc
<b>Table Evidence</b>				
#Samples	482	236	523	256
InternVL3_5-8B	70.6	44.5	68.1	43.8
InternVL3_5-14B	72.0	49.2	70.4	45.7
InternVL3_5-38B	71.7	47.5	72.0	48.4
Qwen3-VL-4B	74.2	49.2	73.2	46.9
Qwen3-VL-8B	75.4	51.7	72.0	48.0
Qwen3-VL-30B-A3B	80.6	62.7	74.6	52.0
o4-mini	85.4	72.5	81.7	65.6
<b>Figure Evidence</b>				
#Samples	265	116	394	180
InternVL3_5-8B	52.2	16.4	54.3	21.7
InternVL3_5-14B	59.7	25.9	56.5	22.8
InternVL3_5-38B	57.9	25.0	62.3	32.8
Qwen3-VL-4B	63.7	32.8	64.9	36.7
Qwen3-VL-8B	65.8	37.1	67.5	38.9
Qwen3-VL-30B-A3B	66.6	38.8	66.3	38.3
o4-mini	78.1	59.5	75.4	53.3

Table 5: Detailed average macro-F1 and pair accuracy scores of the models on our dataset, shown separately for the two types of evidence: tables and figures.

**Table Evidence vs. Figure Evidence.** Table 5 shows the detailed average macro-F1 and pair accuracy scores of the models on our dataset, shown separately for the two types of evidence: tables and figures. As shown in the table, the results of all models on table-based evidence are higher than

those on figure-based evidence. This suggests that samples with table evidence are less challenging for the models, whereas samples with figure evidence are more difficult. For example, on o4-mini, the validation score for table-evidence samples is 85.4 macro-F1, compared to only 78.1 for figure-evidence samples.

#### Figure Evidence Modification Operations.

Based on the previous results, we focus on samples that use figures as evidence. Table 6 shows the models’ average macro-F1 and pair accuracy scores for each type of figure evidence modification.

As discussed in Section 3.4, there are 67 “supported claim only” samples in the figure-based evidence subset. These samples have very low scores, likely because they are difficult even for annotators to verify and create unsupported claims, making them a hard subset. To investigate whether this explains the models’ poorer performance on figure-based evidence compared to table-based evidence, we recalculated results after removing all “supported claim only” samples. The results, shown in Table 7, confirm that reasoning over figures remains more challenging than over tables. To illustrate how “supported claim only” samples affect table-based evidence results, we show detailed scores for different operations in Table 8 in Appendix 5.

Excluding the “supported claim only” subset in Table 6, we observe that different operations vary in difficulty for the models. In general, Others, Graph Swap, and Category Swap are more challenging than Graph Flip.

Model	Validation		Test	
	F1	P-Acc	F1	P-Acc
<b>Graph Flip</b>				
#Samples	40	20	40	20
Qwen3-VL-8B	69.9	45.0	75.9	55.0
Qwen3-VL-30B-A3B	62.3	35.0	72.3	45.0
o4-mini	79.8	60.0	82.9	65.0
<b>Legend Swap</b>				
#Samples	126	63	80	40
Qwen3-VL-8B	65.5	38.1	68.4	42.5
Qwen3-VL-30B-A3B	66.9	39.7	55.3	27.5
o4-mini	81.6	63.5	79.2	57.5
<b>Graph Swap</b>				
#Samples	20	10	18	9
Qwen3-VL-8B	56.3	20.0	45.2	11.1
Qwen3-VL-30B-A3B	71.4	40.0	66.2	33.3
o4-mini	76.8	50.0	68.1	44.4
<b>Category Swap</b>				
#Samples	46	23	26	13
Qwen3-VL-8B	65.0	34.8	70.6	46.2
Qwen3-VL-30B-A3B	66.9	39.1	82.4	61.5
o4-mini	70.3	52.2	72.1	46.2
<b>Others</b>				
#Samples	0	0	196	98
Qwen3-VL-8B	-	-	65.8	35.7
Qwen3-VL-30B-A3B	-	-	66.4	38.8
o4-mini	-	-	74.0	51.0
<b>Supported Claim Only</b>				
#Samples	33		34	
Qwen3-VL-8B	41.1		41.4	
Qwen3-VL-30B-A3B	42.1		43.3	
o4-mini	43.1		41.4	

Table 6: Detailed average macro-F1 and pair accuracy scores of the models using figure-based evidence are shown separately for different types of evidence modification operations.

#### 4.4. Discussion

Considering the human scores reported in Section 3.6, the main results in Section 4.2, and the detailed analyses in Section 4.3, we demonstrate that the figure-based subset of our dataset is challenging for all evaluated models, including o4-mini, as a large gap remains between the best-performing model and the human baseline.

For the table-based subset, although the macro-F1 score is close to the human baseline, there is a notable performance drop when evaluated using pair accuracy, indicating that there is still room for improvement. Moreover, while our table evidence data covers a wide range of table formats, this paper focuses exclusively on the .png format. Conse-

Model	Validation		Test	
	F1	P-Acc	F1	P-Acc
<b>Table Evidence</b>				
#Samples	472	236	512	256
Qwen3-VL-8B	75.2	51.7	72.7	48.0
Qwen3-VL-30B-A3B	80.6	62.7	74.9	52.0
o4-mini	85.5	72.5	82.3	65.6
<b>Figure Evidence</b>				
#Samples	232	116	360	180
Qwen3-VL-8B	65.4	37.1	67.1	38.9
Qwen3-VL-30B-A3B	66.6	38.8	65.9	38.3
o4-mini	78.7	59.5	75.8	53.3

Table 7: Detailed average macro-F1 and pair accuracy scores of the models on our dataset, shown separately for the two types of evidence: tables and figures. We exclude all samples that contain a note indicating the claim is supported only.

quently, substantial research opportunities remain for exploring more diverse table representations within our dataset.

## 5. Conclusion

In this paper, we introduced SciClaimEval, a dataset for scientific claim verification featuring authentic claims, evidence-based negative examples, and diverse data formats. SciClaimEval bridges a key gap between synthetic benchmarks and real-world scientific reasoning. Our evaluation of multiple MLLMs shows that the figure-based subset remains challenging for all models, including o4-mini, with a substantial gap from human performance. In contrast, the table-based subset is more suitable for evaluating open-source MLLMs, as o4-mini achieves near-human performance. Moreover, our table data support multiple formats, providing a valuable resource for further research on scientific paper processing. We hope SciClaimEval will inspire future work on multimodal understanding and the development of more capable and trustworthy scientific reasoning models.

## Limitations

Our research has three main limitations.

First, as described in Section 3.4, we used GPT-5-nano to convert LaTeX and HTML tables into JSON. Human evaluation on 100 random samples revealed 19 cases with major issues and 22 with minor issues, which may affect the quality of the JSON format. We plan to randomly select a subset for human correction in the future.

Second, our dataset contains more table-based than figure-based evidence samples. However, the

o4-mini model’s performance on table-based samples is close to the human baseline, making this subset less challenging.

Third, we only use PNG images for table evidence, leaving other formats unutilized in this work.

## Ethical Statement and Broader Impact

For PeerJ papers, our dataset included 47 papers licensed under CC BY 4.0 and 2 papers licensed under CC BY-NC 4.0. To ensure proper attribution in accordance with these licenses, we added license information, the paper URL, author names, and the paper title to each paper file in JSON format.

For the ML papers, our dataset includes 39 papers licensed under CC BY 4.0, 8 papers licensed under CC BY-NC-SA 4.0, 2 papers licensed under CC BY-SA 4.0, and 2 papers in the public domain. To ensure proper attribution in accordance with these licenses, we added license information, paper URLs, author names, and paper titles to each paper file in JSON format.

For the NLP domain, papers were collected from the ACL Anthology and manually mapped to their arXiv versions when available. Among these, 32 papers are licensed under the arXiv Non-exclusive Distribution License and 4 papers are licensed under CC BY-NC-ND 4.0. Because these licenses do not meet our intended usage requirements, we instead used the ACL versions of these papers, which are licensed under CC BY 4.0. For the remaining papers, we include the arXiv license information (CC BY 4.0: 35 papers; CC BY-NC-SA 4.0: 7 papers; CC BY-SA 4.0: 2 papers).

There are a total of 11 annotators involved in the creation of our dataset. All of them are graduate students or AI/NLP researchers. We do not collect or include any personal or sensitive information in the dataset. Annotators are provided with a detailed guideline during the annotation process. In cases where the guidelines are unclear or ambiguous, they are allowed to provide feedback to the authors of the papers to establish a consistent approach for handling such cases.

## Use of LLMs

We use ChatGPT and GPT-5 to help verify grammar and enhance the quality of our writing. Most of the initial content, however, is authored by us. All suggestions provided by the models are manually reviewed to ensure they accurately convey the intended information. Additionally, we use GitHub Copilot to assist with the coding process.

## Acknowledgements

This work was supported by JSPS KAKENHI Grant Number 24K03231, AID-CNRS NaviTerm project (convention 2022 65 0079 CNRS Occitanie Ouest), and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 554559555.

## Detailed Analysis

Table 8 shows the detailed average macro-F1 and pair accuracy scores of models using table-based evidence, reported separately for different types of evidence modification operations.

Model	Validation		Test	
	F1	P-Acc	F1	P-Acc
<b>Change the Cell Values</b>				
#Samples	366	183	252	126
Qwen3-VL-8B	75.6	53.0	74.7	50.8
Qwen3-VL-30B-A3B	80.7	62.8	75.7	54.8
o4-mini	86.5	73.2	83.4	68.3
<b>Swap Rows or Columns</b>				
#Samples	106	53	222	111
Qwen3-VL-8B	73.6	47.2	71.9	46.8
Qwen3-VL-30B-A3B	80.3	62.3	74.1	49.5
o4-mini	82.0	69.8	82.5	65.8
<b>Alter the Tables</b>				
#Samples	0	0	28	14
Qwen3-VL-8B	-	-	56.0	28.6
Qwen3-VL-30B-A3B	-	-	72.0	42.9
o4-mini	-	-	75.0	50.0
<b>Others</b>				
#Samples	0	0	10	5
Qwen3-VL-8B	-	-	80.0	60.0
Qwen3-VL-30B-A3B	-	-	80.0	60.0
o4-mini	-	-	69.7	40.0
<b>Supported Claim Only</b>				
#Samples	10		11	
Qwen3-VL-8B	44.4		26.7	
Qwen3-VL-30B-A3B	44.4		35.3	
o4-mini	44.4		35.3	

Table 8: Detailed average macro-F1 and pair accuracy scores of the models using table-based evidence are shown separately for different types of evidence modification operations.

## 6. Bibliographical References

- Mubashara Akhtar, Nikesh Subedi, Vivek Gupta, Sahar Tahmasebi, Oana Cocarascu, and Elena Simperl. 2024. [ChartCheck: Explainable fact-checking over real-world chart images](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 13921–13937, Bangkok, Thailand. Association for Computational Linguistics.
- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. [FEVEROUS: Fact extraction and VERification over unstructured and structured information](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. [MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4685–4697, Hong Kong, China. Association for Computational Linguistics.
- Edouard Belval. 2023. [pdf2image](#).
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. [Tabfact: A large-scale dataset for table-based fact verification](#). In *International Conference on Learning Representations*.
- Deyan Ginev. 2024. [ar5iv:04.2024 dataset, an html5 conversion of arxiv.org](#). SIGMathLing – Special Interest Group on Math Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and et al. 2024. [The llama 3 herd of models](#). *arXiv:2407.21783*.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. [A survey on automated fact-checking](#). *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Xanh Ho, Sunisth Kumar, Yun-Ang Wu, Florian Boudin, Atsuhiko Takasu, and Akiko Aizawa. 2025. [Table-text alignment: Explaining claim verification against tables in scientific papers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 2509–2517, Suzhou, China. Association for Computational Linguistics.
- Yichen Jiang, Shikha Bordia, Zheng Zhong, Charles Dognin, Maneesh Singh, and Mohit Bansal. 2020. [HoVer: A dataset for many-hop fact extraction and claim verification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3441–3460, Online.
- Jiho Kim, Sungjin Park, Yeonsu Kwon, Yohan Jo, James Thorne, and Edward Choi. 2023. [FactKG: Fact verification via reasoning on knowledge graphs](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16190–16206, Toronto, Canada. Association for Computational Linguistics.
- Neema Kotonya and Francesca Toni. 2020. [Explainable automated fact-checking for public health claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7740–7754, Online. Association for Computational Linguistics.
- Ashish Kulkarni. 2023. [wkhtmltopdf](#).
- Yash Kumar Lal, Manikanta Bandham, Mohammad Saqib Hasan, Apoorva Kashi, Mahnaz Koupaee, and Niranjana Balasubramanian. 2025. [MuSciClaims: Multimodal scientific claim verification](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 3285–3307, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. [Llava-onevision: Easy visual task transfer](#). *arXiv:2408.03326*.
- Xinyuan Lu, Liangming Pan, Qian Liu, Preslav Nakov, and Min-Yen Kan. 2023. [SCITAB: A challenging benchmark for compositional reasoning and claim verification on scientific tables](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7787–7813, Singapore. Association for Computational Linguistics.
- Meta-AI. 2025. [“the llama 4 herd: The beginning of a new era of natively multimodal intelligence”](#).
- OpenAI. 2025. [Addendum to openai o3 and o4-mini system card: Openai o3 operator](#).

- 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.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or fiction: Verifying scientific claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics.
- David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang, and Hannaneh Hajishirzi. 2022. [SciFact-open: Towards open-domain scientific claim verification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4719–4734, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chengye Wang, Yifei Shen, Zexi Kuang, Arman Cohan, and Yilun Zhao. 2025a. [SciVer: Evaluating foundation models for multimodal scientific claim verification](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8562–8579, Vienna, Austria. Association for Computational Linguistics.
- Nancy X. R. Wang, Diwakar Mahajan, Marina Danilevsky, and Sara Rosenthal. 2021. [SemEval-2021 task 9: Fact verification and evidence finding for tabular data in scientific documents \(SEM-TAB-FACTS\)](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 317–326, Online. Association for Computational Linguistics.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, Guanzhou Chen, Zichen Ding, Changyao Tian, Zhenyu Wu, Jingjing Xie, Zehao Li, Bowen Yang, Yuchen Duan, Xuehui Wang, Zhi Hou, Haoran Hao, Tianyi Zhang, Songze Li, Xiangyu Zhao, Haodong Duan, Nianchen Deng, Bin Fu, Yinan He, Yi Wang, Conghui He, Botian Shi, Junjun He, Yingtong Xiong, Han Lv, Lijun Wu, Wenqi Shao, Kaipeng Zhang, Huipeng Deng, Biqing Qi, Jiaye Ge, Qipeng Guo, Wenwei Zhang, Songyang Zhang, Maosong Cao, Junyao Lin, Kexian Tang, Jianfei Gao, Haiyan Huang, Yuzhe Gu, Chengqi Lyu, Huanze Tang, Rui Wang, Haijun Lv, Wanli Ouyang, Limin Wang, Min Dou, Xizhou Zhu, Tong Lu, Dahua Lin, Jifeng Dai, Weijie Su, Bowen Zhou, Kai Chen, Yu Qiao, Wenhai Wang, and Gen Luo. 2025b. [InternV3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency](#). *arXiv:2508.18265*.
- William Yang Wang. 2017. [“liar, liar pants on fire”: A new benchmark dataset for fake news detection](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 422–426, Vancouver, Canada.
- Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. 2004. [Image quality assessment: from error visibility to structural similarity](#). *IEEE Transactions on Image Processing*, 13(4):600–612.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chu-jie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025a. [Qwen3 technical report](#). *arXiv:2505.09388*.
- Bohao Yang, Yingji Zhang, Dong Liu, André Freitas, and Chenghua Lin. 2025b. [Does table source matter? benchmarking and improving multimodal scientific table understanding and reasoning](#). *arXiv:2501.13042*.
- Yuji Zhang, Qingyun Wang, Cheng Qian, Jiateng Liu, Chenkai Sun, Denghui Zhang, Tarek Abdelzaher, Chengxiang Zhai, Preslav Nakov, and Heng Ji. 2025. [Atomic reasoning for scientific table claim verification](#). *arXiv:2506.06972*.