

# Do Large Multimodal Models Solve Caption Generation for Scientific Figures? Lessons Learned from SciCAP Challenge 2023

Ting-Yao ‘Edward’ Hsu<sup>1</sup> Yi-Li Hsu<sup>2</sup> Shaurya Rohatgi<sup>3</sup> Chieh-Yang Huang<sup>4</sup>  
Ho Yin Sam Ng<sup>1</sup> Ryan Rossi<sup>5</sup> Sungchul Kim<sup>5</sup> Tong Yu<sup>5</sup> Lun-Wei Ku<sup>6</sup>  
Clyde Lee Giles<sup>1</sup> Ting-Hao ‘Kenneth’ Huang<sup>1</sup>

<sup>1</sup>Pennsylvania State University, USA <sup>2</sup>National Tsing Hua University, Taiwan

<sup>3</sup>AllSci, USA <sup>4</sup>MetaMetrics Inc., USA <sup>5</sup>Adobe Research, USA

<sup>6</sup>Institute of Information Science, Academia Sinica, Taiwan

<sup>1</sup>{txh357, sam.ng, clg20, txh710}@psu.edu <sup>3</sup>srohatgi@allsci.com

<sup>4</sup>cyhuang@lexile.com <sup>5</sup>{ryrossi, sukim, tyu}@adobe.com

<sup>2,6</sup>{yili.hsu, lwku}@iis.sinica.edu.tw

## Abstract

Since the SciCAP dataset’s launch in 2021, the research community has made significant progress in generating captions for scientific figures in scholarly articles. In 2023, the first SciCAP Challenge took place, inviting global teams to use an expanded SciCAP dataset to develop models for captioning diverse figure types across various academic fields. At the same time, text generation models advanced quickly, with many powerful pre-trained large multimodal models (LMMs) emerging that showed impressive capabilities in various vision-and-language tasks. This paper presents an overview of the first SciCAP Challenge and details the performance of various models on its data, capturing a snapshot of the field’s state. We found that professional editors overwhelmingly preferred figure captions generated by GPT-4V over those from all other models and even the original captions written by authors. Following this key finding, we conducted detailed analyses to answer this question: **Have advanced LMMs solved the task of generating captions for scientific figures?**

## 1 Introduction and Background

Scientists use figures like bar charts, pie charts, or scatter plots to convey key findings in scholarly articles. However, the texts accompanying these figures, the figure captions, are often overlooked by the authors and do not receive the needed attention when being composed. Even though many studies have shown the role of captions in enhancing readers’ comprehension and recall of the messages conveyed by figures (Nugent, 1983; Large et al., 1995; Bransford, 1979; Hegarty

and Just, 1993), poorly written captions are, unfortunately, common (Huang et al., 2023).

In response, Hsu et al. (2021) launched SciCAP, a large-scale collection of 133,543 single-panel line charts and their captions extracted from arXiv papers, aiming to fuel the creation of new models that generate high-quality captions for scientific figures. Over the past three years, SciCAP has advanced researchers’ understanding of scientific figure captions and driven the progress of technologies generating them: SciCAP confirmed that many low-quality captions exist in scholarly articles, as over half of the figure captions in arXiv cs.CL papers were rated ‘‘unhelpful’’ by NLP Ph.D. students (Huang et al., 2023); SciCAP revealed that producing figure captions in scholarly articles is a generative task heavily reliant on the texts within the articles (Yang et al., 2023; Li and Tajbakhsh, 2023; Horawalavithana et al., 2023)—this task relies so much on the paper content that it can be more effectively tackled through text summarization, summarizing all paragraphs mentioning the figures (*e.g.*, ‘‘Figure 3 shows...’’), rather than as a vision-to-language task (Huang et al., 2023); SciCAP elevated the quality of generated captions to a level where, in instances where author-created captions were poorly crafted, readers found the generated captions more helpful (Huang et al., 2023), thereby offering practical assistance in caption writing (Hsu et al., 2024). In 2023, the first SciCAP Challenge took place. With an expanded SciCAP of 476,389 single-panel figures from 8 domains and 5 figure types, the challenge invited teams worldwide to develop caption-generation models for diverse figures (Figure 1).

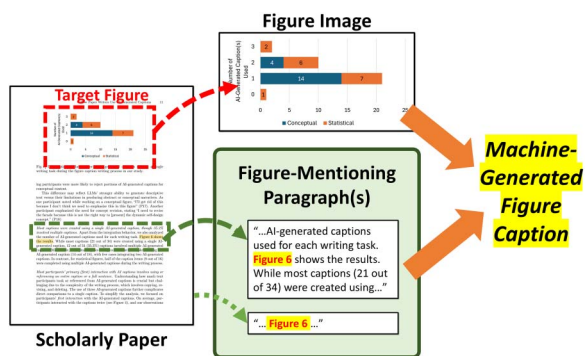


Figure 1: In the SciCAP Challenge, models generate captions based on the figure and the figure-mentioning paragraph. [Paper source: Ng et al. (2025)].

Meanwhile, in the past two years, the landscape of text generation has rapidly evolved: ChatGPT (GPT-3.5) was initially released to the public in November 2022; its more powerful successor, GPT-4, was released in March 2023; and GPT-4V, the successor to GPT-4, which can take images and texts as input, became publicly available in September 2023. Numerous open-source pre-trained Large Language Models (LLMs) and Large Multimodal Models (LMMs), such as OPT (Zhang et al., 2022), LLaMA-2 (Touvron et al., 2023), Mistral (Jiang et al., 2023), Gemma (Team et al., 2024), LLaVA (Liu et al., 2023), BLIP-2 (Li et al., 2023), mPLUG-Owl2 (Ye et al., 2023), and MiniGPT-4 (Zhu et al., 2023), have been made available. All these models have shown impressive progress in understanding charts and tables within documents (Yue et al., 2023). Now is a good time to take a step back and critically assess the collective advancements made in generating high-quality, useful captions for scientific figures. Specifically, we aim to answer this question: **Have the impressive large multimodal models solved the challenge of generating good captions for scientific figures?**

This paper first overviews the 2023 SciCAP Challenge, including its data, procedure, and winning teams and their models, which represents the status quo of the scientific figure caption generation task (Section 3). Next, we ran GPT-4V—unavailable at the time of the challenge—and other open LMMs, such as UniChart (Masry et al., 2023), on SciCAP Challenge’s data and, through automatic and human evaluations, compared their performances to all models that participated in the challenge (Section 4). The automatic evaluation revealed that the linguis-

tic features of captions differ by domain and figure type, resulting in a model’s scores varying across these categories. More interestingly, in the human evaluation, three professional editors with expertise in technical academic writing **unanimously preferred captions generated by GPT-4V over those from all other models, including the original captions written by the authors.** Driven by this key finding, we investigated why editors strongly prefer GPT-4V (Section 6.3), comparing their views with those of Ph.D. and undergrad students (Section 5), to directly address whether LMMs have solved scientific figure caption generation. We conclude by identifying unresolved challenges and suggesting future research directions (Section 6).

## 2 Related Work

In addition to the works directly related to SciCAP (Section 1), the complex nature of figure narrating—which demands an understanding of both vision and text, along with intricate domain-specific contexts—has made it an intellectually intriguing challenge, inspiring various parallel projects focused on text generation for figures. FigJAM leveraged metadata and a combined static and dynamic dictionary to generate better caption units (Qian et al., 2021). FigCAP developed innovative attention mechanisms and employed reinforcement learning for sequence-level training to enhance caption generation performance (Chen et al., 2020). Large-scale chart-to-text benchmark datasets have also been created to facilitate data-to-text techniques for chart summarization (Kantharaj et al., 2022; Obeid and Hoque, 2020). Tarsi et al. (2024) presented two datasets for scientific multimodal learning: SciOL, a large corpus covering various sciences, and MuLMS-Img, a high-quality annotated materials science dataset. Xia et al. (2024) introduced ChartX, a comprehensive evaluation set with 18 chart types, 7 tasks, 22 topics, and developed ChartVLM, a model tailored for multi-modal tasks requiring interpretable patterns, like chart reasoning. Unichart developed specialized pretraining tasks for charts, focusing on visual element extraction and reasoning skills (Masry et al., 2023). ChartSumm introduced a comprehensive benchmark dataset with 84,363 charts, metadata, and descriptions across various topics and types for creating concise and detailed summaries (Rahman

et al., 2023). While inspiring, most efforts did not aim to produce real-world captions from scholarly articles. To show this, we included UniChart’s output in our study.

### 3 The First SciCAP Challenge (2023)

The first SciCAP Challenge was held at the 5th Workshop on Closing the Loop Between Vision and Language during ICCV 2023. The organizers expanded the SciCAP dataset to produce a challenge-specific version which was divided into test, validation, and training sets, containing 476,389 single-panel arXiv figures. Teams had three months to develop solutions and submit model outputs for all test set figures. This section details the challenge.

#### 3.1 Challenge Dataset

The SciCAP challenge dataset comprises 476,389 single-panel figures from arXiv papers<sup>1</sup> in 8 domains: (i) Computer Science (cs), (ii) Economics (econ), (iii) Electrical Engineering and Systems Science (eess), (iv) Mathematics (math), (v) Physics (physics), (vi) Quantitative Biology (q-bio), (vii) Quantitative Finance (q-fin), and (viii) Statistics (stat). It covers 5 figure types, as denoted by the figure type classifier, FigureSeer (Siegel et al., 2016): (i) Node Diagram, (ii) Equation, (iii) Graph Plot, (iv) Scatterplot, and (v) Bar Chart. This dataset, with 476,389 figures, is more than 3 times the size of the original SciCAP dataset, which contains 133,543 figures. The figures were randomly sampled from arXiv papers published between 2010 and 2020, with a maximum of four figures selected per paper. While the organizers had access to the full collection of over 440,000 arXiv papers from this period, they intentionally limited the dataset size to encourage participation for teams with limited computational resources. Appendix A details the dataset preparation process.

Concerns about the quality of arXiv papers are understandable. However, Huang et al. (2023) manually verified 399 figures from the original SciCAP test set and found that 81.2% (324/399) originated from papers published at academic conferences, with 51.9% (207/399) appearing in the ACL Anthology, IEEE, or ACM.<sup>2</sup> These findings

<sup>1</sup>The organizing team also supplied 26,868 figures from ACL papers as extra training data, which were not included in the data mentioned in this section.

<sup>2</sup>These results were reported in the paper’s Appendix B.

suggest that the dataset is both representative and of reasonable quality.

**Data Split.** The dataset was divided as follows: training set (70%, 333,472 figures), validation set (10%, 47,639 figures), public test set (10%, 47,639 figures), and hidden test set (10%, 47,639 figures). Teams could submit results for the public test set throughout the three-month submission period to refine their systems. Access to the hidden test set was granted in the final two weeks, during which teams could submit their hidden test set results. The winning teams were determined based on performance on the hidden test set.

**Selecting the “Quality Subset” from the Hidden Test Set.** To address issues with unreliable data skewing reference-based evaluations, the SciCAP Challenge organizers curated a “Quality Subset” of the hidden test set, composed only of high-quality captions. Starting with a random sample of 1,000 figures from the hidden test set, they applied Hsu et al.’s (2023) approach to have GPT-4 assign each caption a quality score between 1 and 6. Captions scoring below 5 were discarded, resulting in 769 captions. The organizers then manually evaluated these captions, ultimately selecting 460 captions for the “Quality Subset.” This subset was used to determine the Quality Winner.

**Data Provided.** For each figure in the training, validation, and public test set, the following data items were provided: (i) figure image, (ii) figure caption, (iii) figure type, (iv) OCR (textual information extracted from the image), and (v) the paragraphs that mentioned the figure (e.g., “As shown in Figure 5, ...”). All of these were extracted from the original arXiv paper. For each figure in the hidden test set, only the image and the paragraphs, but not the caption, were provided.

**Full-Scale SciCAP Data Release.** Following the SciCAP Challenge 2023 and 2024 (Kim et al., 2025), and in response to growing interest in large-scale figure caption data, the organizing team decided to release the full SciCAP dataset, including all figure captions extracted from arXiv papers published between 2010 and 2020.<sup>3</sup> Future SciCAP Challenges will use newly published arXiv papers as test and hidden test data.

<sup>3</sup>SciCAP Full Data Release: <https://huggingface.co/datasets/CrowdAILab/Scicap-All>.

### 3.2 Challenge Procedure

The SciCap Challenge was hosted on EvalAI (Yadav et al., 2019). It had two phases: the **Test Phase** and the **Challenge Phase**. During the Test Phase—which began on May 29, 2023, and lasted approximately 2.5 months—each team used the provided training set, validation set, and public test set to build and test their models. Each team could upload their predictions for the public test set to Eval.AI, which automatically calculates the scores using a series of evaluation metrics. Then, in the Challenge Phase, which took place in the final 2 weeks before the submission deadline (August 15–31, 2023), the hidden test set was released. Teams could then submit their results for the hidden test set. Each team could submit up to five times per day during the Challenge Phase. To qualify for the winners, the team must have submitted their prediction outputs for the hidden test set and provided a two to four-page technical report.

#### Rules for Using External Data and Models.

Teams were allowed to use LLMs like GPT-4 or LLaMA and any external data for their captioning systems. However, using the original captions from the hidden test set was prohibited to maintain fairness. Although we did not provide gold captions in the hidden test set, arXiv papers are openly accessible, making it impossible to fully prevent potential data leakage. Teams caught using this gold data were disqualified. Each team must clearly describe their methodology in the technical report to enable replication.

**Evaluation and Winning Criteria.** The challenge selected two winners: the **Leaderboard Winner** and the **Quality Winner**.<sup>4</sup> The challenge’s leaderboard reported seven automatic scores that were commonly used in prior works (Huang et al., 2023; Hsu et al., 2021), using author-written captions as references: BLEU-4, ROUGE-1, ROUGE-2, ROUGE-L, and the normalized versions of ROUGE scores. Note that ROUGE’s F1 scores were used. The **ROUGE-2-Normalized** score was chosen as the deciding metric for determining the winners, as it addresses the issue of standard ROUGE F1

<sup>4</sup>The challenge was set to award two distinct winning teams. If one team topped both evaluations, the runner-up in the Quality Winner track becomes the winner for that category.

Team	BLEU-4	ROUGE-1		ROUGE-2		ROUGE-L	
		Score	Norm	Score	Norm	Score	Norm
<b>Hidden Set</b>							
NJUST	.082	<b>.406</b>	<b>2.331</b>	<b>.250</b>	<b>4.489</b>	<b>.378</b>	<b>2.617</b>
USTC	.043	.297	1.613	.145	2.418	.256	1.700
Pegasus	<b>.099</b>	<u>.363</u>	1.880	.220	<u>3.462</u>	<u>.325</u>	<u>2.086</u>
Pegasus*	<u>.088</u>	.340	1.832	.200	3.313	.304	2.001
<b>Quality Subset</b>							
NJUST	.070	<u>.372</u>	<b>2.043</b>	<u>.226</u>	<b>3.828</b>	<b>.338</b>	<b>2.260</b>
USTC	.049	.307	1.612	.157	2.503	.260	1.680
Pegasus	<b>.118</b>	<b>.384</b>	<u>1.887</u>	<b>.240</b>	<u>3.507</u>	<b>.338</b>	<u>2.090</u>
Pegasus*	<u>.096</u>	.358	1.825	.214	3.293	.312	1.981

Table 1: Leaderboard scores of the two winning teams and the two baselines. The **best** and **second-best** results are highlighted. Note that Pegasus is fine-tuned on the Scicap Challenge dataset and Pegasus\* is fine-tuned on graphplots only (Huang et al., 2023).

favoring longer texts (Huang et al., 2023; Sun et al., 2019). The Leaderboard Winner achieved the highest ROUGE-2-Normalized score on the entire hidden test set, and the Quality Winner achieved the highest ROUGE-2-Normalized score on the specially curated Quality Subset of the hidden test set (Section 3.1). The Quality Winner track was inspired by Huang et al. (2023), highlighting that a model performing well on a full test set may struggle with higher-quality captions.

### 3.3 Challenge Results and Winning Teams

Team NJUST-KMG (Chao et al., 2023) scored highest on both public and hidden test sets, with Team USTC-IAT-United (Yu et al., 2023) as the runner-up in the Quality Winner track. According to the rules, NJUST-KMG was named the Leaderboard Winner, and USTC-IAT-United was selected as the Quality Winner. Table 1 shows the evaluation scores of the winning teams on the entire Hidden Test Set and the Quality Subset, including scores from two baseline models the SciCAP Challenge reported in the leaderboard.

**Leaderboard Winner: NJUST-KMG (Chao et al., 2023).** The team employed the PP-OCRv3 model from PaddleOCR (Du et al., 2020) to extract precise information from images. They used LLaMA-2-7B to refine the paragraphs with summarization. This refined dataset was then used to train a PegasusX model (Phang et al., 2022) to generate candidate summaries. Another BRIO model (Liu et al., 2022) was trained with these

summaries, using contrastive learning to optimize ROUGE-2-normalized performance.

**Quality Winner: USTC-IAT-United (Yu et al., 2023).** The team first enriched the dataset by concatenating BLIP-v2 (Li et al., 2023) captions, paragraphs, and mentions. Then a Pegasus model (Zhang et al., 2020) was fine-tuned on this enriched dataset. Although the team tried different summarization models, the fine-tuned Pegasus was shown to be the most effective one.

## 4 In-Depth Evaluation of Models' Performance

We conducted a series of analyses on each model's performance, including comparisons with advanced LMMs that were unavailable during the challenge. Study 1 (Section 4.1) and Study 2 (Section 4.2) assessed performance using the challenge's hidden test set, with Study 1 focusing on automatic evaluation results and Study 2 on human evaluation results. To verify the generalizability of our findings, Study 3 (Section 4.4) repeated these evaluations on newer arXiv papers.

### 4.1 Models Included in the Evaluation

This analysis compared the models of the two winning teams with three other sets of models:

**Text Summarization Model: Pegasus.** Prior studies showed that generating captions for scientific figures in scholarly articles can be effectively tackled by treating it as a text summarization task, *i.e.*, summarizing all paragraphs mentioning the target figure into its caption (Huang et al., 2023). We fine-tuned Pegasus (Zhang et al., 2020), a well-known text summarization model, with the SciCap Challenge dataset's training set in our performance comparison.

**Large Multimodal Model (LMM): GPT-4V.<sup>5</sup>** We evaluated GPT-4V for figure captioning in two settings: (i) using **figure-mentioning paragraphs and images** as input, and (ii) using **only figure images**. We used GPT-4V through OpenAI's API. Appendix B describes the details.

<sup>5</sup>We used *gpt-4-vision-preview* in the main study. However, due to its deprecation on June 17, 2024, we employed *gpt-4-0125-preview* in the Section 4.1.1 and Section 4.4.

**Open LMM: UniChart (Masry et al., 2023).** We fine-tuned UniChart using the SciCAP Challenge data. The training and inference details are provided in Appendix C.

### 4.1.1 Why Exclude Text-Only LMMs?

In our main study, we evaluated only LMMs that can take images as input, excluding text-only conditions. This decision followed our observation that combining text and images results in better captions. A human evaluation confirmed this performance difference: We randomly sampled 200 figures from the arXiv Challenge Dataset (Section 3) and generated two captions for each figure using GPT-4V: one caption based solely on the textual context (*i.e.*, paragraphs referencing the figure), and another that also incorporated the figure images. The same three experts who performed human evaluation in Study 2 (Section 4.3) ranked these captions. Experts ranked image+text captions higher than text-only captions in 82.5%, 66%, and 56% of cases, respectively. This indicates that captions with images (mean = 1.32; lower scores denote better performance) significantly outperformed text-only captions (mean=1.68).

### 4.2 Study 1: Automatic Evaluation Results

We evaluated our models using BLEU-4, ROUGE-1, ROUGE-2, and ROUGE-L metrics (Papineni et al., 2002; Lin, 2004), calculating ROUGE scores with the `rouge-score` tool on all-lowercase, stemmed text (Google Research, 2022). Following Huang et al. (2023), we also used normalized ROUGE scores because ROUGE can be affected by text length, with longer texts typically scoring higher (Sun et al., 2019). The normalization factor was computed on the hidden test set. Table 2 and Table 3 show model performance across domains and figure types.

**Models Maintained Consistent Performance Rankings Across Categories.** Figure 2 and Figure 3 show the ROUGE-2-Normalized scores for each model across various domains and figure types. The Leaderboard Winner, NJUST, consistently scored the highest in every category, followed by PEGASUS in second place and USTC in third. Scores for both LLMs and LMMs were consistently lower. Model score rankings were consistent across categories, but score ranges varied. The domains of `cs` and `eess` saw the highest

Model	Domain	Token Length	BLEU-4	ROUGE-1		ROUGE-2		ROUGE-L	
				Score	Norm	Score	Norm	Score	Norm
NJUST	physics	14.4	<u>.062</u>	.323	1.856	<u>.195</u>	3.494	<u>.287</u>	1.986
	cs	8.3	<u>.078</u>	<u>.373</u>	2.950	<u>.230</u>	6.317	<u>.349</u>	3.143
	math	9.5	<u>.070</u>	<u>.339</u>	2.447	<u>.205</u>	4.975	<u>.317</u>	2.638
	q-bio	11.4	<u>.062</u>	.307	1.990	<u>.190</u>	3.991	<u>.279</u>	2.129
	stat	9.9	<u>.068</u>	.334	2.348	<u>.201</u>	4.691	<u>.310</u>	2.520
	q-fin	10.3	<u>.053</u>	.319	2.191	<u>.193</u>	4.385	<u>.293</u>	2.347
	eess	9.3	<u>.100</u>	<u>.414</u>	3.031	<u>.260</u>	6.431	<u>.386</u>	3.260
	econ	10.0	<u>.067</u>	<u>.361</u>	2.534	<u>.214</u>	4.985	<u>.331</u>	2.690
USTC	physics	21.6	.050	.319	1.584	.156	2.316	.268	1.665
	cs	10.9	.048	.314	2.095	.155	3.378	.284	2.212
	math	12.8	.043	.308	1.871	.150	2.892	.277	2.008
	q-bio	15.9	.042	.291	1.603	.140	2.373	.250	1.674
	stat	13.1	.048	.306	1.838	.149	2.829	.271	1.940
	q-fin	14.1	.044	.295	1.712	.150	2.732	.261	1.818
	eess	11.4	.061	.347	2.245	.178	3.734	.315	2.394
	econ	13.2	.048	.321	1.926	.156	2.953	.291	2.078
Pegasus	physics	24.4	<b>.109</b>	<b>.384</b>	1.859	<b>.232</b>	3.312	<b>.334</b>	2.045
	cs	12.4	<b>.118</b>	<b>.394</b>	2.448	<b>.244</b>	4.843	<b>.367</b>	2.699
	math	14.3	<b>.113</b>	<b>.391</b>	2.259	<b>.243</b>	4.389	<b>.364</b>	2.525
	q-bio	19.8	<b>.106</b>	<b>.367</b>	1.870	<b>.229</b>	3.525	<b>.329</b>	2.086
	stat	16.3	<b>.122</b>	<b>.384</b>	2.092	<b>.242</b>	4.074	<b>.352</b>	2.338
	q-fin	16.2	<b>.102</b>	<b>.361</b>	1.970	<b>.228</b>	3.837	<b>.329</b>	2.188
	eess	12.8	<b>.133</b>	<b>.428</b>	2.605	<b>.269</b>	5.186	<b>.397</b>	2.870
	econ	17.8	<b>.152</b>	<b>.439</b>	2.314	<b>.292</b>	4.695	<b>.409</b>	2.651
Unichart (Fine-tuned)	physics	20.3	.005	.174	0.879	.049	0.738	.148	0.932
	cs	8.1	.007	.184	1.481	.065	1.830	.168	1.539
	math	13.1	.004	.163	0.981	.045	0.851	.149	1.068
	q-bio	13.2	.008	.185	1.107	.058	1.092	.159	1.137
	stat	11.1	.005	.176	1.159	.056	1.194	.156	1.207
	q-fin	13.6	.006	.177	1.041	.055	1.012	.157	1.106
	eess	9.3	.013	.219	1.609	.081	2.007	.200	1.697
	econ	12.5	.003	.154	0.950	.050	0.990	.141	1.033
GPT-4V (Image)	physics	20.7	.004	.203	1.019	.049	0.742	.157	0.984
	cs	17.5	.005	.219	1.159	.059	0.960	.179	1.165
	math	17.6	.002	.168	0.889	.034	0.548	.139	0.906
	q-bio	18.2	.005	.213	1.114	.059	0.930	.168	1.081
	stat	17.5	.003	.214	1.135	.053	0.864	.170	1.106
	q-fin	18.0	.003	.191	0.999	.046	0.737	.154	0.996
	eess	18.1	.005	.234	1.226	.062	0.991	.191	1.233
	econ	18.8	.004	.212	1.100	.067	1.055	.183	1.171
GPT-4V (Image+Paragraph)	physics	41.3	.025	<u>.352</u>	1.529	.132	1.582	.258	1.480
	cs	28.0	.017	.324	1.537	.124	1.704	.257	1.559
	math	30.9	.020	.313	1.458	.119	1.582	.249	1.489
	q-bio	35.0	.020	<u>.333</u>	1.509	.122	1.564	.245	1.439
	stat	31.4	.018	<u>.336</u>	1.559	.125	1.662	.256	1.527
	q-fin	31.0	.024	<u>.328</u>	1.527	.127	1.698	.257	1.535
	eess	29.6	.022	.335	1.573	.130	1.766	.268	1.613
	econ	31.3	.022	.330	1.530	.130	1.733	.266	1.588

Table 2: Results breakdown by domain for each model, highlighting performance metrics such as BLEU-4, ROUGE scores, and token length. We show the models included in the human evaluation process. The **highest** and second highest values across models are highlighted.

Model	Domain	Token Length	BLEU-4	ROUGE-1		ROUGE-2		ROUGE-L	
				Score	Norm	Score	Norm	Score	Norm
NJUST	Node Diagram	6.5	<u>.079</u>	<u>.359</u>	3.452	<u>.231</u>	8.369	<u>.342</u>	3.657
	Equation	8.6	<u>.068</u>	<u>.341</u>	2.650	<u>.212</u>	5.690	<u>.318</u>	2.826
	Graph Plot	13.0	<u>.068</u>	<u>.342</u>	2.066	<u>.206</u>	3.950	<u>.309</u>	2.223
	Scatterplot	13.5	<u>.057</u>	<u>.321</u>	1.901	<u>.190</u>	3.544	<u>.284</u>	2.010
	Bar Chart	10.0	<u>.080</u>	<u>.380</u>	2.666	<u>.230</u>	5.355	<u>.350</u>	2.850
USTC	Node Diagram	9.2	.050	.304	2.252	.154	3.869	.281	2.391
	Equation	13.4	.043	.295	1.750	.141	2.642	.263	1.868
	Graph Plot	18.5	.051	.321	1.669	.158	2.499	.275	1.769
	Scatterplot	19.9	.042	.309	1.568	.144	2.215	.257	1.630
	Bar Chart	13.2	.053	.338	2.025	.165	3.131	.298	2.130
Pegasus	Node Diagram	11.6	<b>.140</b>	<b>.406</b>	2.611	<b>.269</b>	5.595	<b>.386</b>	2.919
	Equation	15.7	<b>.113</b>	<b>.379</b>	2.095	<b>.233</b>	3.990	<b>.347</b>	2.334
	Graph Plot	20.7	<b>.110</b>	<b>.387</b>	1.946	<b>.234</b>	3.531	<b>.344</b>	2.156
	Scatterplot	22.8	<b>.098</b>	<b>.372</b>	1.819	<b>.220</b>	3.189	<b>.322</b>	1.989
	Bar Chart	14.6	<b>.117</b>	<b>.402</b>	2.289	<b>.241</b>	4.282	<b>.366</b>	2.518
Unichart (Fine-tuned)	Node Diagram	5.7	.002	.112	1.202	.027	1.134	.107	1.256
	Equation	9.1	.001	.097	0.723	.019	0.472	.090	0.768
	Graph Plot	17.1	.007	.198	1.061	.064	1.056	.172	1.131
	Scatterplot	21.1	.004	.180	0.899	.050	0.751	.154	0.962
	Bar Chart	9.9	.008	.222	1.567	.084	1.975	.201	1.645
GPT-4V (Image)	Node Diagram	16.7	.001	.170	0.918	.034	0.562	.141	0.931
	Equation	18.0	.001	.149	0.781	.027	0.423	.122	0.791
	Graph Plot	19.9	.005	.222	1.125	.059	0.904	.174	1.102
	Scatterplot	20.8	.004	.202	1.012	.049	0.735	.155	0.970
	Bar Chart	16.6	.005	.238	1.286	.069	1.148	.197	1.301
GPT-4V (Image+Paragraph)	Node Diagram	29.0	.014	.297	1.402	.112	1.525	.239	1.443
	Equation	35.9	.010	.255	1.146	.087	1.099	.200	1.174
	Graph Plot	36.5	.025	<u>.353</u>	1.583	.135	1.697	.266	1.552
	Scatterplot	40.2	.024	<u>.354</u>	1.547	.132	1.604	.259	1.490
	Bar Chart	27.7	.019	.351	1.667	.133	1.830	.277	1.678

Table 3: Results breakdown by the figure type for each model, highlighting performance metrics such as BLEU-4, ROUGE scores, and token length. We show the models included in the human evaluation. The **highest** and second highest values across models are highlighted.

scores, and Node Diagrams emerged as the figure type with the highest scores.

We speculate that differences in automatic scores across categories stem from the linguistic characteristics of captions rather than the amount of training data available. `physics`, the domain with the lowest scores in Figure 2, was the most common, and graph plots, the most frequent figure type, also scored low in Figure 3.

### 4.3 Study 2: Human Evaluation Results By Professional Editors

**Expert Judges.** One of the primary goals of the SciCAP dataset and challenge is to produce captions that help human readers. To assess this,

we recruited three professional editors through UpWork (upwork.com), specializing in academic articles in the technical field and all native American English speakers. Their qualifications include: one with over ten years of editing experience and a Ph.D. in Comparative Literature, and two from STEM fields—one in Theoretical Astrophysics and another in Neuroscience—both with a significant track record in editing, proofreading, and publishing academic papers. They spent between 30 and 60 minutes evaluating 10 figures, with their rates ranging from \$50 to \$60 per hour.

**Three Varying Length Constraints.** These editors conducted a three-part human evaluation.




Figure 2: ROUGE-2 normalized scores of each model across eight arXiv domains, highlighting similar trends and demonstrating the generalizability of the caption generation approaches.




Figure 3: ROUGE-2 scores by model across five figure types, showing similar trends.

Caption length is known to be a key factor influencing human judgments of quality, as longer captions are often perceived as more informative by readers (Hartley, 2003; Gelman et al., 2002).<sup>6</sup> To assess human-perceived model performance under varying length constraints, we conducted human evaluations for captions generated under three different settings, ranging from no length constraints to the strictest length constraints:

1. **Generation with no length constraints (all arXiv domains, Figure 4A):** For the first part, we selected 100 figures from the hidden test set and asked each editor to rank 6 captions for each figure: 1 author-written and 5 machine-generated by models used in Section 4.2 (including the two teams’ outputs), excluding GPT-4V with image-only input due to its poor quality. They used a drag-and-drop interface (Figure 9 in Appendix D), similar to that of Hsu et al., to rank the captions based on the criterion: “When I read the paper, the caption can help me understand the message that the figure tries to convey.” This represented the least length-constrained setting. For additional context, when no length constraints were specified in prompts, the average length

<sup>6</sup>Note that we did not deliberately control the generation lengths in Study 1, as the length biases in automatic metrics were taken care of by metrics normalization.

of GPT-4V’s captions (34.9 words) was shorter than that of human-written captions (42.9 words).

2. **Generation with 25-word length constraints (cs Papers, Figure 4B):** The second part followed the same procedure as the first, but focused specifically on 100 figures from `cs` domain papers. We selected the `cs` domain because the average length of human-written captions in `cs` arXiv papers is 25.52 words, shorter than GPT-4V’s average caption length (33.6 words) when no length constraints are given. This choice helped mitigate the potential effect of longer human-written captions in the first condition. To align GPT-4V’s captions with the average length in the `cs` domain, we modified the prompt to limit caption length to 25 words. Consequently, the average length of GPT-4V’s captions was 25.79 words, closely matching the 25-word target.
3. **Generation with length no longer than human captions (all arXiv domains, Figure 4C):** In the third part, we followed the same procedure as in the first condition to sample 100 figures from papers of all arXiv domains. We modified GPT-4V’s prompt to restrict caption length to be no longer than the corresponding human-written captions from arXiv papers, establishing the strictest



Figure 4: Rankings of generated captions by all models in Study 2 across three evaluation conditions (A, B, C) and three experts. Models were ranked from 1 (best, green) to 6 (worst, red). GPT-4V (Image+Paragraph) consistently outperformed other models, including humans, across varying length constraints: none (A), a 25-word limit (B), and a strict limit matching human-written caption length (C).

length constraints among all three settings. In this setting, generated captions were expected to always be shorter or match the length of human-written captions.

**GPT-4V Captions Were Preferred Over Author-Written Ones and All Other Models.** The ranking results are shown in Figure 4. Captions generated by GPT-4V, using both paragraphs and images, were overwhelmingly preferred by all experts in every setting, surpassing even author-written captions. Although author-written captions ranked slightly lower than GPT-4V, they still outperformed all other models, remaining reliable when considering the proportion of outputs ranked in the top two. Among the remaining models, the text-only summarization model, Pegasus, consistently outperformed Unichart and the two team models. Unichart, which was designed for chart understanding and reasoning rather than text generation, received the lowest scores in human evaluation. It is worth emphasizing that extensive evidence shows that automatic evaluation metrics, like BLEU, do not align well with human judgment (Dhingra et al., 2019). Therefore, in this paper, we place greater

emphasis on human evaluation over automatic evaluation results. We also analyzed the experts’ free-text feedback to understand their rankings, as detailed in Section 6.3.

**Length Constraints Did Not Affect GPT-4V’s Superiority.** The human evaluation results show that, although human readers tend to favor longer captions, imposing stricter length constraints did not diminish the strong preference for GPT-4’s captions. Each expert displayed individual preferences, but overall, they favored GPT-4’s captions in most cases. Notably, GPT-4 roughly adhered to length constraints: In Condition 1 (all arXiv domains), the average lengths were 41.69 for human-written captions and 33.6 for GPT-4. In Condition 2 (CS domain only), the averages were 25.52 for humans and 25.79 for GPT-4. In Condition 3 (all arXiv domains with human-written length constrained), the human average was 39.68, while GPT-4’s average was 44.59.

**Inter-Annotator Agreements.** We measured inter-annotator agreements between experts using Kendall’s Tau (for ranking all six items) and Cohen’s Kappa (for selecting the top 1 or top 2 items). The Kendall’s Tau values for Condition

Model	Domain	Token Length	BLEU-4	ROUGE-1		ROUGE-2		ROUGE-L	
				Score	Norm	Score	Norm	Score	Norm
Pegasus	physics	39.1	<b>.096</b>	<b>.378</b>	1.651	<b>.208</b>	2.804	<b>.311</b>	1.864
	cs	18.9	<b>.054</b>	<b>.316</b>	1.613	<b>.157</b>	2.674	<b>.275</b>	1.796
	math	22.1	<b>.055</b>	<b>.321</b>	1.563	<b>.168</b>	2.684	<b>.282</b>	1.791
	q-bio	29.6	<b>.067</b>	<u>.317</u>	1.475	<b>.162</b>	2.415	<b>.272</b>	1.690
	stat	25.0	<b>.061</b>	<u>.339</u>	1.622	<b>.178</b>	2.757	<b>.289</b>	1.819
	q-fin	21.4	<b>.051</b>	<b>.327</b>	1.606	<b>.166</b>	2.684	<b>.285</b>	1.815
	eess	19.6	<b>.072</b>	<b>.344</b>	1.736	<b>.182</b>	3.056	<b>.303</b>	1.964
	econ	22.3	<b>.027</b>	<b>.274</b>	1.331	<b>.124</b>	1.976	<b>.236</b>	1.499
GPT-4V (Image+Paragraph)	physics	77.8	<u>.032</u>	<u>.375</u>	1.461	<u>.137</u>	1.483	<u>.267</u>	1.518
	cs	53.3	<u>.018</u>	<u>.307</u>	1.225	<u>.110</u>	1.301	<u>.240</u>	1.362
	math	60.8	<u>.018</u>	<u>.297</u>	1.170	<u>.112</u>	1.283	<u>.237</u>	1.340
	q-bio	59.4	<u>.019</u>	<b>.331</b>	1.303	<u>.117</u>	1.350	<u>.250</u>	1.415
	stat	61.3	<u>.025</u>	<b>.345</b>	1.360	<u>.132</u>	1.508	<u>.264</u>	1.491
	q-fin	52.5	<u>.021</u>	<u>.318</u>	1.276	<u>.121</u>	1.443	<u>.251</u>	1.430
	eess	59.1	<u>.017</u>	<u>.290</u>	1.145	<u>.103</u>	1.186	<u>.229</u>	1.292
	econ	57.9	<u>.016</u>	<u>.257</u>	1.012	<u>.099</u>	1.140	<u>.211</u>	1.190
Unichart (Fine-tuned)	physics	25.2	.004	.171	0.816	.045	0.695	.143	0.903
	cs	9.1	.004	.176	1.289	.059	1.590	.159	1.354
	math	15.5	.002	.152	0.834	.040	0.746	.136	0.930
	q-bio	11.7	.004	.165	1.037	.052	1.152	.146	1.101
	stat	13.7	.003	.175	1.018	.056	1.120	.154	1.100
	q-fin	11.8	.003	.185	1.161	.067	1.479	.166	1.253
	eess	11.9	.007	.201	1.250	.076	1.663	.183	1.373
	econ	14.5	.008	.148	0.839	.047	0.906	.133	0.928

Table 4: Results breakdown by domain for each model, highlighting performance metrics such as BLEU-4, ROUGE scores, and token length. We show the models included in the human evaluation process. We test the fine-tuned model in Table 2 on unseen data. The **highest** and second highest values across models are highlighted.

A, B, and C across expert pairs (1–2, 2–3, 3–1) were (0.71, 0.68, 0.64), (0.66, 0.53, 0.59), and (0.6, 0.58, 0.56), respectively, indicating a strong positive correlation between experts. Given that our caption ranking study prioritizes the selection of top-ranked captions—since users are more likely to adopt these as drafts or read them—we also evaluated agreement using Cohen’s Kappa, treating the ranking task as a binary classification task: captions ranked first were labeled as 1, and all others as 0 (Kappa@1). Similarly, Kappa@2 classified the top two captions. For Kappa@1, the values for Conditions A, B, and C were (0.74, 0.64, 0.70), (0.78, 0.81, 0.72), and (0.63, 0.59, 0.53), respectively. For Kappa@2, the corresponding values were (0.72, 0.78, 0.72), (0.66, 0.60, 0.52), and (0.61, 0.55, 0.58). These results indicate substantial agreement between experts.

#### 4.4 Study 3: Are Results Generalizable to Papers Posted After the Cut-off Date?

A common concern with large pre-trained models is data contamination (Balloccu et al., 2024), meaning that test data may overlap with data in the pretraining set. To address this, we repeated the experiments on **entirely unseen data for GPT-4V**, creating a new dataset from papers published after the model’s data cut-off date, January 2024.

**Experiment Setups.** After filtering out tables and other figure types, this dataset initially contained 46,543 figures from 57,678 valid papers published between January and March 2024. Of these, 37,146 figures had associated paragraphs. We further refined the dataset by removing duplicate figure indices within the same paper, resulting in 36,841 figures. Finally, after excluding figures

Model	Domain	Token Length	BLEU-4	ROUGE-1		ROUGE-2		ROUGE-L	
				Score	Norm	Score	Norm	Score	Norm
Pegasus	Node Diagram	18.3	<b>.056</b>	<b>.303</b>	1.564	<b>.160</b>	2.770	<b>.270</b>	<b>1.780</b>
	Equation	22.3	<b>.050</b>	<b>.303</b>	1.473	<b>.146</b>	2.335	<b>.263</b>	1.669
	Graph Plot	30.2	<b>.080</b>	<b>.357</b>	1.655	<b>.190</b>	2.812	<b>.301</b>	1.867
	Scatterplot	35.1	<b>.080</b>	<u>.359</u>	1.609	<b>.186</b>	2.612	<b>.294</b>	1.790
	Bar Chart	21.0	<b>.056</b>	<u>.330</u>	1.632	<b>.161</b>	2.635	<b>.284</b>	1.815
GPT-4V (Image+Paragraph)	Node Diagram	56.8	<u>.012</u>	<u>.279</u>	1.102	<u>.093</u>	1.083	<u>.218</u>	1.228
	Equation	49.7	<u>.014</u>	<u>.279</u>	1.138	<u>.092</u>	1.123	<u>.218</u>	1.256
	Graph Plot	68.9	<u>.027</u>	<u>.346</u>	1.354	<u>.128</u>	1.423	<u>.257</u>	1.457
	Scatterplot	70.8	<u>.028</u>	<b>.366</b>	1.433	<u>.133</u>	1.471	<u>.263</u>	1.491
	Bar Chart	51.9	<u>.023</u>	<b>.338</b>	1.360	<u>.126</u>	1.511	<u>.263</u>	1.500
Unichart (Fine-tuned)	Node Diagram	6.8	.001	.120	1.085	.031	1.088	.111	1.128
	Equation	9.9	.000	.109	0.751	.023	0.572	.101	0.819
	Graph Plot	19.7	.006	.189	0.954	.060	1.003	.164	1.063
	Scatterplot	22.4	.002	.171	0.829	.044	0.699	.145	0.920
	Bar Chart	9.9	.004	.202	1.400	.072	1.807	.181	1.473

Table 5: Results are broken down by domain for each model, highlighting performance metrics such as BLEU-4, ROUGE scores, and token length. We include the models used in the human evaluation process and test the fine-tuned model in Table 3 on unseen data. The **highest** and second highest values across models are highlighted.




Figure 5: Rankings of captions generated by different models on an unseen dataset of arXiv papers posted after GPT-4V’s data cut-off date, evaluated by three experts. The procedure replicated Study 2.

with invalid captions, the dataset contained 36,740 figures, which is similar in size to the challenge’s hidden test set. We used *gpt-4-0125-preview* here to generate figure captions (See Footnote 5).

Similar to the challenge task, we conducted both automatic evaluation (Table 4 and Table 5) and human evaluation (Figure 5). Note that two teams’ models were excluded, as this study was not part of the SciCAP Challenge.

**The Results from Studies 1 and 2 Generalize Well to Unseen arXiv Figures.** As discussed in Study 2, this work places a greater emphasis on human evaluation. The human evaluation results from Study 3 (Figure 5) showed a similar trend to those in Study 2 (Figure 4), with GPT-4V’s captions consistently favored by all experts. Additionally, Study 3’s automatic evaluation results (Table 4 and Table 5) displayed trends similar to those in Study 1 (Table 2 and

Table 3), where text-summarization models received higher scores. Overall, the findings from Studies 1 and 2 appear to generalize to figures in unseen arXiv papers.

## 5 Additional Study: Do Paper Readers Agree with the Editors’ Judgment?

Study 2 (Section 4.3) shows that professional editors’ perspective toward machine-generated captions, raising an intriguing question: Do paper *readers* evaluate caption quality in the same way as professional editors? This distinction is important because prior studies have often relied on general readers, such as graduate students, rather than professional editors for human evaluations (Petsiuk et al., 2022; Kasai et al., 2022). Understanding any systematic differences between readers and editors is essential to inform the broader research community.




Figure 6: Ph.D. students’ ranking results from Hsu et al. (2023). Note that lower ranks mean better performance.




Figure 7: Undergraduate students’ ratings on helpfulness from Hsu et al. (2023). Note that higher ratings mean better performance.

To answer this question, we obtained the data for the study conducted by Hsu et al. (2023), where Ph.D. students in the relevant fields ranked figure captions generated by different models, and undergraduate students rated the captions’ helpfulness. The models Hsu et al. (2023) used differed from ours but similarly employed a fine-tuned text summarization model, Pegasus, as a baseline. Their study included three versions of captions produced by Pegasus: one fine-tuned with the entire training set of the SciCAP dataset [Pegasus (Paragraph+OCR)], one fine-tuned using only captions exceeding 30 words in length [Pegasus (Paragraph+OCR+Better)],<sup>7</sup> and the other fine-tuned using only the figure image’s OCR [Pegasus (OCR)]. They also included author-written captions in their comparison, providing a basis for comparing their results with ours. Interestingly, their study focused on three *cs* subdomains of arXiv: natural language processing (NLP, *cs.CL*), human-computer interaction (HCI, *cs.HC*), and computer vision (CV, *cs.CV*). This section describes the insight gained from comparing their results, as shown in Figures 6 and 7, with ours.

### Students Agreed with Editors that the *Basic Text-Summarization Model Did Not Outperform Paper Authors.* Ph.D. students (Figure 6)

<sup>7</sup>It was called “Better” because previous studies indicated that longer captions enhanced reader comprehension (Hartley, 2003; Gelman et al., 2002).

and undergraduate students (Figure 7) both preferred author-written captions to those generated by the Pegasus (Paragraph+OCR) model, which was fine-tuned on the entire dataset rather than just long captions. These preferences aligned with the editors’ judgments (Section 4.3), where human-authored captions were chosen more often over Pegasus-generated ones, even though our text-summarization model was fine-tuned on a much larger dataset than Hsu et al.’s (2023).

**User Groups Do Not Always Agree with Each Other.** While the overall trend in quality assessment among models was consistent between Ph.D. students and undergraduates—such as the high rating of human-written captions—agreement was not universal. Undergraduates found captions from the Pegasus model trained on longer captions more useful than human-written ones (Figure 7), a view not shared by Ph.D. students (Figure 6). This discrepancy shows diverse user needs (Gkatzia et al., 2014), beyond what editors’ ratings capture.

## 6 Discussion

### 6.1 Is the Problem Solved?

Our human evaluations show that GPT-4V often generates captions that are seen as better than those written by authors (Section 4.3). However, this does not fully solve the challenge of generating captions for scientific figures. We have reached an important milestone, but there are new challenges ahead:

- First, **improving caption quality further.** Being better than human-written captions does not necessarily mean the captions are of high quality; it might simply mean that the human-written captions are bad.
- Second, **evaluation is still hard.** Despite GPT-4V’s poor automatic scores (Section 4.2), it outperformed other models in evaluations by editors; efforts to use a higher-quality subset for better evaluation did not significantly differ in scores from the full test set (Table 1). These highlight the need for more research into evaluation methods, particularly as the quality of machine-generated content surpasses that of human-generated content. Future research should focus on detail, coherence, and

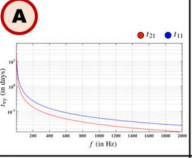
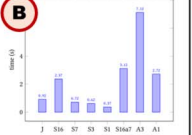
Figure Image	Caption Written By Paper Authors	Caption Generated By GPT-4V (Image + Paragraphs)
	<b>Figure 10:</b> $t_{21}$ characteristics: $t_{21}$ and $t_{11}$ are plotted as a function of neutron star's rotational frequency $f$ . We have set $v_c^0 = 0.1c$ , $N^2 = 10^{-4}$ , $E = 10^{-10}$ , $\epsilon = 10^{-4}$ , $d_{\text{source}} = 1 \text{ kpc}$ , $L = 10^4 \text{ m}$ , $g = 10^{12} \text{ m/sec}^2$ , $\rho_0 = 10^{17} \text{ Kg/m}^3$ , $\partial v_c = 0$ .	<b>Figure:</b> Graph depicting the variation of mode timescales $t_{21}$ (red) and $t_{11}$ (blue) with neutron star rotational frequency $f$ . The timescales, spanning several orders of magnitude, are crucial for determining the minimum detectable strain in gravitational wave searches, with a notable dip in mid-frequency range due to the sign change in $V_{\nu\gamma}$ .
	<b>Figure 2:</b> AoT compilation time of the Eigen benchmark. See the text for a description of the configurations presented.	<b>Figure 3:</b> Performance comparison of JIT-compiled, AoT specialization, and AoT generic code for matrix operations with sizes 3, 7, and 16, showing JIT-compiled code matches or exceeds AoT specialization performance and significantly outperforms AoT generic as size increases.

Figure 8: Two examples where experts favored GPT-4V’s captions for providing sufficient details and highlighting key takeaway messages, two key factors identified in our comment analysis. [Figure sources: (A) Singh (2017), (B) Finkel et al. (2019)].

factual accuracy. Developing a reliable automatic evaluation method could also enhance the quality of datasets used to develop new models.

- Finally, **personalization for different users**. Different user groups have different information needs. Another challenge is to create models that can customize captions to meet the diverse preferences of various users.

## 6.2 Generalizability of Caption Generation Models

Our findings suggest caption generation models are versatile across domains and figure types. Although there are some variations and domain-specific details, large multimodal models pre-trained on extensive data can likely generate usable captions for many domains and figure types. We also tested several hypotheses to explain the variations in automatic scores across domains, such as the presence of non-English characters in captions or length of captions, but found no conclusive evidence to support them.

## 6.3 What Makes GPT-4’s Captions Better?

To understand why GPT-4’s generated captions were often rated higher than author-written captions, we manually analyzed the comments provided by three experts during the ranking process. Although experts were instructed to provide comments, we did not strictly enforce this for every figure. In total, we collected 1,123 comments corresponding to 500 figures. The first author of this paper manually coded these comments using open coding practices to identify the reasons behind GPT-4V’s superior ratings. Two key themes emerged: (i) providing **sufficient details**

and (ii) highlighting the figure’s **takeaway messages**. Among the 1,123 comments analyzed, 277 (24.7%) indicated that captions were rated higher or lower based on the presence or absence of rich details, while 177 (15.8%) noted that including or omitting takeaway messages influenced their ratings. These findings align with prior research showing that detailed captions with clear takeaway messages receive higher ratings (Hsu et al., 2023). Because humans tend to prefer longer captions, we further analyzed the 207 comments from Study 2’s C condition, where GPT-4’s captions were restricted to the same length as the authors’. Even under this constraint, the two key factors remained prominent: 49 (23.7%) comments mentioned details, and 50 (24.2%) comments cited takeaway messages. Figure 8 shows two representative examples where experts stated that GPT-4V’s captions were better due to these factors.

We also analyzed the experts’ comments to identify the most common errors in AI-generated captions, particularly those from models other than GPT-4. These errors fell into three main categories: The first category is **incorrect visual information**, where visual details in captions, such as color-number relationships, are inaccurately described. The second category involves **linguistic errors**, including typos, missing punctuation, or minor grammatical mistakes. The final category is **factual errors unrelated to visuals**, such as incorrect numerical values. In addition to outright errors, we also examined cases where captions, though not strictly incorrect, were problematic. One common issue was **incomplete captions**, where truncation resulted in phrases like “Delivery ratio vs.” or “Plot of F2 versus ab.” Another problem was **repetitive or overly**

**generic captions**, where models produced identical or indistinct descriptions, such as ‘‘IoT Hub Workflow,’’ reflecting a lack of specificity.

## 7 Conclusion and Future Work

This paper presents the first SciCAP Challenge results, highlighting GPT-4V’s superiority in generating scientific figure captions, as professional editors overwhelmingly preferred figure captions generated by GPT-4V over those from all other models and even the original captions written by authors. Despite this achievement, we acknowledge that while an important milestone has been reached in caption generation, challenges such as improving quality, evaluation, and customization remain unresolved.

## 8 Limitations

This work has several limitations. First, our human evaluation involved only a small number of experts—three professional editors. While this scale is not uncommon in NLP research due to the high cost of expert evaluation, it may not sufficiently capture the diverse perspectives and interpretations of a broader group of users or experts. Second, the comparison made in Section 5 between our results and those from Hsu et al. (2023) was not fully direct due to differences in data and models, which we made clear in the paper. While we believe the comparison is valid at a high level, it may not be broadly generalizable. Third, our analysis spans all domains present on arXiv, but arXiv does not encompass all potential academic domains, such as biology and medicine, which are primarily published on PubMed. The characteristics of figures and captions in these areas may differ significantly from those in arXiv, potentially limiting the applicability of our conclusions outside of arXiv without additional verification. Finally, our best-performing model, GPT-4V, is proprietary, with no public access to its architecture or training data. This restriction, common in NLP research, means our findings may inherit limitations related to the closed nature of the model, including potential data contamination issues, and lacks transparency for in-depth analysis.

## 9 Ethics Statement

We recognize that employing LLMs or LMMs to produce texts for end users inherently carries risks, including the dissemination of inaccurate

or misleading information. In scholarly contexts, particularly when using generated captions to enhance figure comprehension, such inaccuracies could mislead readers. We have ensured to inform participants of these risks during our user studies.

## Acknowledgments

We sincerely appreciate the anonymous reviewers and action editors of *Transactions of the Association for Computational Linguistics* for their valuable feedback. We also thank the human experts for their support in our evaluation efforts. Additionally, we express our gratitude to all the teams that participated in SciCAP Challenge 2023 and to the organizers of the 5th Workshop on Closing the Loop Between Vision and Language (CLVL) at ICCV 2023 for hosting the SciCAP Challenge 2023. This research was partially supported by the Alfred P. Sloan Foundation (grant number: 2024-22721).

## References

- Simone Balloccu, Patrícia Schmidtová, Mateusz Lango, and Ondrej Dusek. 2024. Leak, cheat, repeat: Data contamination and evaluation malpractices in closed-source LLMs. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 67–93, St. Julian’s, Malta. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.eacl-long.5>
- John Bransford. 1979. *Human Cognition: Learning, Understanding, and Remembering*. Wadsworth Publishing.
- Vitaly Bulgakov. 2016. Iterative aggregation method for solving principal component analysis problems. *arXiv preprint arXiv:1602.08800*.
- Dian Chao, Xin Song, Shupeng Zhong, Boyuan Wang, Xiangyu Wu, Chen Zhu, and Yang Yang. 2023. The solution for the ICCV 2023 1st scientific figure captioning challenge. *arXiv preprint*.
- Charles Chen, Ruiyi Zhang, Eunye Koh, Sungchul Kim, Scott Cohen, and Ryan Rossi. 2020. Figure captioning with relation maps for reasoning. In *Proceedings of the IEEE/CVF*

- Winter Conference on Applications of Computer Vision*, pages 1537–1545. <https://doi.org/10.1109/WACV45572.2020.9093592>
- Christopher Clark and Santosh Divvala. 2016. Pdffigures 2.0: Mining figures from research papers. In *Proceedings of the 16th ACM/IEEE-CS on Joint Conference on Digital Libraries*, pages 143–152. <https://doi.org/10.1145/2910896.2910904>
- Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. Handling divergent reference texts when evaluating table-to-text generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4884–4895, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1483>
- Yuning Du, Chenxia Li, Ruoyu Guo, Xiaoting Yin, Weiwei Liu, Jun Zhou, Yifan Bai, Zilin Yu, Yehua Yang, Qingqing Dang, and Haoshuang Wang. 2020. PP-OCR: A practical ultra lightweight OCR system. *arXiv preprint arXiv:2009.09941*.
- Hal Finkel, David Poliakoff, Jean-Sylvain Camier, and David F. Richards. 2019. Clangjit: Enhancing c++ with just-in-time compilation. In *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, pages 82–95. IEEE. <https://doi.org/10.1109/P3HPC49587.2019.00013>
- Andrew Gelman, Cristian Pasarica, and Rahul Dodhia. 2002. Let’s practice what we preach: Turning tables into graphs. *The American Statistician*, 56(2):121–130. <https://doi.org/10.1198/000313002317572790>
- Dimitra Gkatzia, Helen Hastie, and Oliver Lemon. 2014. Finding middle ground? Multi-objective natural language generation from time-series data. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers*, pages 210–214, Gothenburg, Sweden. Association for Computational Linguistics. <https://doi.org/10.3115/v1/E14-4041>
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Google Research. 2022. Python rouge implementation. <https://github.com/google-research/google-research/tree/master/rouge>
- James Hartley. 2003. Single authors are not alone: Colleagues often help. *Journal of Scholarly Publishing*, 34(2):108–113. <https://doi.org/10.3138/jsp.34.2.108>
- Mary Hegarty and Marcel-Adam Just. 1993. Constructing mental models of machines from text and diagrams. *Journal of Memory and Language*, 32(6):717–742. <https://doi.org/10.1006/jmla.1993.1036>
- Sameera Horawalavithana, Sai Munikoti, Ian Stewart, and Henry Kvinge. 2023. Scitune: Aligning large language models with scientific multimodal instructions. *arXiv preprint arXiv:2307.01139*. <https://doi.org/10.18653/v1/2024.nlp4science-1.7>
- Ting-Yao Hsu, C. Lee Giles, and Ting-Hao Huang. 2021. SciCap: Generating captions for scientific figures. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3258–3264.
- Ting-Yao Hsu, Chieh-Yang Huang, Shih-Hong Huang, Ryan Rossi, Sungchul Kim, Tong Yu, C. Lee Giles, and Ting-Hao K. Huang. 2024. SciCapenter: Supporting caption composition for scientific figures with machine-generated captions and ratings. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems*.
- Ting-Yao Hsu, Chieh-Yang Huang, Ryan Rossi, Sungchul Kim, C. Giles, and Ting-Hao Huang. 2023. GPT-4 as an effective zero-shot evaluator for scientific figure captions. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5464–5474. <https://doi.org/10.18653/v1/2023.findings-emnlp.363>
- Chieh-Yang Huang, Ting-Yao Hsu, Ryan Rossi, Ani Nenkova, Sungchul Kim, Gromit Yeuk-Yin Chan, Eunyee Koh, Clyde Lee Giles, and Ting-Hao ‘Kenneth’ Huang. 2023. Summaries as captions: Generating figure captions for scientific documents with automated

- text summarization. <https://doi.org/10.18653/v1/2023.inlg-main.6>
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Shankar Kantharaj, Rixie Tiffany Ko Leong, Xiang Lin, Ahmed Masry, Megh Thakkar, Enamul Hoque, and Shafiq Joty. 2022. Chart-to-text: A large-scale benchmark for chart summarization. *arXiv preprint arXiv:2203.06486*.
- Zeba Karishma, Shaurya Rohatgi, Kavya Shrinivas Puranik, Jian Wu, and C. Lee Giles. 2023. ACL-fig: A dataset for scientific figure classification. *arXiv preprint arXiv:2301.12293*.
- Jungo Kasai, Keisuke Sakaguchi, Lavinia Dunagan, Jacob Morrison, Ronan Le Bras, Yejin Choi, and Noah A. Smith. 2022. Transparent human evaluation for image captioning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3464–3478, Seattle, United States. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.naacl-main.254>
- Jaeyoung Kim, Jongho Lee, Hong-Jun Choi, Ting-Yao Hsu, Chieh-Yang Huang, Sungchul Kim, Ryan Rossi, Tong Yu, Clyde Lee Giles, Ting-Hao’Kenneth’ Huang, and Sungchul Choi. 2025. Multi-LLM collaborative caption generation in scientific documents. *arXiv preprint arXiv:2501.02552*.
- Rodney Michael Kinney, Chloe Anastasiades, Russell Authur, Iz Beltagy, Jonathan Bragg, Alexandra Buraczynski, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Arman Cohan, Miles Crawford, Doug Downey, Jason Dunkelberger, Oren Etzioni, Rob Evans, Sergey Feldman, Joseph Gorney, David W. Graham, F. Q. Hu, Regan Huff, Daniel King, Sebastian Kohlmeier, Bailey Kuehl, Michael Langan, Daniel Lin, Haokun Liu, Kyle Lo, Jaron Lochner, Kelsey MacMillan, Tyler Murray, Christopher Newell, Smita R. Rao, Shaurya Rohatgi, Paul L. Sayre, Zejiang Shen, Amanpreet Singh, Luca Soldaini, Shivashankar Subramanian, A. Tanaka, Alex D. Wade, Linda M. Wagner, Lucy Lu Wang, Christopher Wilhelm, Caroline Wu, Jiangjiang Yang, Angele Zamarron, Madeleine van Zuylen, and Daniel S. Weld. 2023. The semantic scholar open data platform. *ArXiv*, abs/2301.10140.
- Andrew Large, Jamshid Beheshti, Alain Breuleux, and Andre Renaud. 1995. Multi-media and comprehension: The relationship among text, animation, and captions. *Journal of the American Society for Information Science*, 46(5):340–347. [https://doi.org/10.1002/\(SICI\)1097-4571\(199506\)46:5<340::AID-ASI5>3.0.CO;2-S](https://doi.org/10.1002/(SICI)1097-4571(199506)46:5<340::AID-ASI5>3.0.CO;2-S)
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models.
- Shengzhi Li and Nima Tajbakhsh. 2023. Sci-GraphQA: A large-scale synthetic multi-turn question-answering dataset for scientific graphs. *arXiv preprint arXiv:2308.03349*.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81. Barcelona, Spain. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. BRIO: Bringing order to abstractive summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2890–2903, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.207>
- Ahmed Masry, Parsa Kavehzadeh, Xuan Long Do, Enamul Hoque, and Shafiq Joty. 2023. UniChart: A universal vision-language pre-trained model for chart comprehension and

- reasoning. <https://doi.org/10.18653/v1/2023.emnlp-main.906>
- Ho Yin Sam Ng, Ting-Yao Hsu, Jiyoo Min, Sungchul Kim, Ryan A. Rossi, Tong Yu, Hyunggu Jung, and Ting-Hao Kenneth Huang. 2025. Understanding how paper writers use AI-generated captions in figure caption writing. In *2nd AI4Research Workshop: Towards a Knowledge-grounded Scientific Research Lifecycle*.
- Gwen C. Nugent. 1983. Deaf students' learning from captioned instruction: The relationship between the visual and caption display. *The Journal of Special Education*, 17(2):227–234. <https://doi.org/10.1177/002246698301700214>
- Jason Obeid and Enamul Hoque. 2020. Chart-to-text: Generating natural language descriptions for charts by adapting the transformer model. *arXiv preprint arXiv:2010.09142*. <https://doi.org/10.18653/v1/2020.inlg-1.20>
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318. <https://doi.org/10.3115/1073083.1073135>
- Vitali Petsiuk, Alexander E. Siemenn, Saisamrit Surbehera, Zad Chin, Keith Tyser, Gregory Hunter, Arvind Raghavan, Yann Hicke, Bryan A. Plummer, Ori Kerret, Tonio Buonassisi, Kate Saenko, Armando Solar-Lezama, and Iddo Drori. 2022. Human evaluation of text-to-image models on a multi-task benchmark.
- Jason Phang, Yao Zhao, and Peter J. Liu. 2022. Investigating efficiently extending transformers for long input summarization. <https://doi.org/10.18653/v1/2023.emnlp-main.240>
- Xin Qian, Eunyee Koh, Fan Du, Sungchul Kim, Joel Chan, Ryan A. Rossi, Sana Malik, and Tak Yeon Lee. 2021. Generating accurate caption units for figure captioning. In *Proceedings of the Web Conference 2021*, pages 2792–2804. <https://doi.org/10.1145/3442381.3449923>
- Raian Rahman, Rizvi Hasan, Abdullah Al Farhad, Md Tahmid Rahman Laskar, Md Hamjajul Ashmafee, and Abu Raihan Mostofa Kamal. 2023. ChartSumm: A comprehensive benchmark for automatic chart summarization of long and short summaries. *arXiv preprint arXiv:2304.13620*.
- Shaurya Rohatgi, Yanxia Qin, Benjamin Aw, Niranjana Unnithan, and Min-Yen Kan. 2023. The ACL OCL corpus: Advancing open science in computational linguistics. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10348–10361, Singapore. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.640>
- Noah Siegel, Zachary Horvitz, Roie Levin, Santosh Divvala, and Ali Farhadi. 2016. Figureseer: Parsing result-figures in research papers. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pages 664–680. Springer. [https://doi.org/10.1007/978-3-319-46478-7\\_41](https://doi.org/10.1007/978-3-319-46478-7_41)
- Avneet Singh. 2017. Gravitational wave transient signal emission via ekman pumping in neutron stars during post-glitch relaxation phase. *Physical Review D*, 95(2):024022. <https://doi.org/10.1103/PhysRevD.95.024022>
- Simeng Sun, Ori Shapira, Ido Dagan, and Ani Nenkova. 2019. How to compare summarizers without target length? Pitfalls, solutions and re-examination of the neural summarization literature. In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 21–29, Minneapolis, Minnesota. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W19-2303>
- Tim Tarsi, Heike Adel, Jan Hendrik Metzen, Dan Zhang, Matteo Finco, and Annemarie Friedrich. 2024. SciOL and MuLMS-Img: Introducing a large-scale multimodal scientific dataset and models for image-text tasks in the scientific domain. In *Proceedings of the IEEE/CVF Winter Conference on*

*Applications of Computer Vision (WACV)*, pages 4560–4571. <https://doi.org/10.1109/WACV57701.2024.00450>

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Satoshi Tsutsui and David Crandall. 2017. A data driven approach for compound figure separation using convolutional neural networks. In *The IAPR International Conference on Document Analysis and Recognition (ICDAR)*. <https://doi.org/10.1109/ICDAR.2017.93>
- Renqiu Xia, Bo Zhang, Hancheng Ye, Xiangchao Yan, Qi Liu, Hongbin Zhou, Zijun Chen, Min Dou, Botian Shi, Junchi Yan, and Yu Qiao. 2024. ChartX & chartVLM: A versatile benchmark and foundation model for complicated chart reasoning. *arXiv preprint arXiv:2402.12185*.
- Deshraj Yadav, Rishabh Jain, Harsh Agrawal, Prithvijit Chattopadhyay, Taranjeet Singh, Akash Jain, Shiv Baran Singh, Stefan Lee, and Dhruv Batra. 2019. EvaLAI: Towards better evaluation systems for AI agents.
- Zhishen Yang, Raj Dabre, Hideki Tanaka, and Naoaki Okazaki. 2023. Scicap+: A knowledge augmented dataset to study the challenges of scientific figure captioning. *arXiv preprint arXiv:2306.03491*.
- Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Haowei Liu, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. 2023. mPLUG-Owl2: Revolutionizing multi-modal large language model with modality collaboration. *arXiv preprint arXiv:2311.04257*.
- Jun Yu, Xiaohua Qi, Biao Zhang, and Qiang Ling. 2023. A solution to the 1st scientific figure captioning (scicap) challenge. *arXiv preprint*.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. 2023. MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI. *arXiv preprint arXiv:2311.16502*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. OPT: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. MiniGPT-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.

## A SciCAP Challenge Data Preparation

The SciCAP Challenge dataset was an expansion of the original SciCAP dataset, incorporating papers from all 8 primary arXiv domains published between 2010 and 2020, totaling over 440,000 papers. Using PDFFigures 2.0 (Clark and Divvala, 2016), 2.62 million figure-caption

pairs were extracted from these papers. The dataset also included paragraphs referencing these figures, extracted using GROBID,<sup>8</sup> and identified figure-mentioning sentences (e.g., “Figure 4 shows...”) through regular expressions. The dataset featured OCR text from images obtained via Tesseract-OCR,<sup>9</sup> enhancing the data with legends and labels directly from figures. A figure-separator tool (Tsutsui and Crandall, 2017) was used to spot compound figures; FigureSeer (Siegel et al., 2016) was used to classify figure types, excluding those categorized as “Table” and “Others”. Finally, to accommodate teams with limited resources and ensure a balanced dataset, we limited the selection to four figures per paper, with figures from the same paper allocated to the same dataset split.

**Additional Figures from ACL Papers.** Additionally, the challenge dataset included 26,868 figures collected from ACL Anthology papers (Rohatgi et al., 2023) through ACL-Fig (Karishma et al., 2023) as additional training data. In the ACL collection, we removed all the papers that were already published at arXiv. For this filtering, we relied on the paper clustering done by Semantic Scholar Academic Graph (Kinney et al., 2023).

## B Prompts Used

In this section, we provide the prompt we used in Section 4.1. [PARAGRAPHS] and [IMG]

are placeholders for figure-mentioning paragraphs and the encoded images.

**Prompt without paragraph.** “[IMG] Please write a short and concise caption for the figure.”

**Prompt with paragraph.** “[IMG] Paragraphs: [PARAGRAPHS]. Above are a figure and referred paragraphs about the figure. Please write a short and concise caption for the figure.”

## C Model Details

We describe the model training details and the decoding configuration used in Section 4.1.

**Training Details for Open-Source LMMs.** We fine-tune Unichart<sup>10</sup> checkpoint from HuggingFace for 40,000 steps, using batch size = 16, learning rate = 5e-5 with a linear decay scheduler, warmup steps = 500, weight decay = 0.01. We evaluate every 1000 steps, and the checkpoint with the highest ROUGE-2 score on validation set is kept and used to predict final result.

**Decoding Details for Open-Source LMMs.** For generation, captions were decoded using the beam sampling strategy, with beam size = 4, top-k = 50, and maximum length = 100.

## D Interface for Human Evaluation

Figure 9 shows the interface used for human evaluation.

<sup>8</sup>GROBID: <https://github.com/kermitt2/grobid>.

<sup>9</sup>Tesseract-OCR: <https://github.com/tesseract-ocr/tesseract>.

<sup>10</sup>We used [ahmed-masry/unichart-base-960](https://github.com/ahmed-masry/unichart-base-960).

Please drag the boxes below, **ranking them** based on the following criteria:

**When I read the paper, the caption can help me understand the message that the figure tries to convey.**

Some of these captions were generated by computers, which might contain some errors.

- When the caption is obviously wrong, and almost all the information in it is incorrect, it is a bad caption.
- When the caption is generally fine but has a few obvious factual errors, judge based on whether or not the caption helps understand the figure despite the errors. (In practice, readers will likely be informed that the caption was **auto-generated**.)
- Each **box index** refers to the left-side **figure index**. **Re-order** the boxes based on your understanding of caption quality, **the prior the better**.
- Please **disregard the figure index** mentioned in the caption during your evaluation.

Strongly Agree / Best Caption

Caption 1

Caption 2

Caption 3

Caption 4

Caption 5

Caption 6

Strongly Disagree / Worst Caption

Progress: 1/1

Previous Figure
Next Figure

Notes for this figure or caption (optional):

Comments here


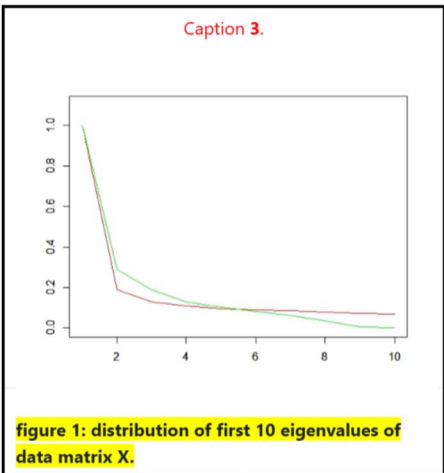

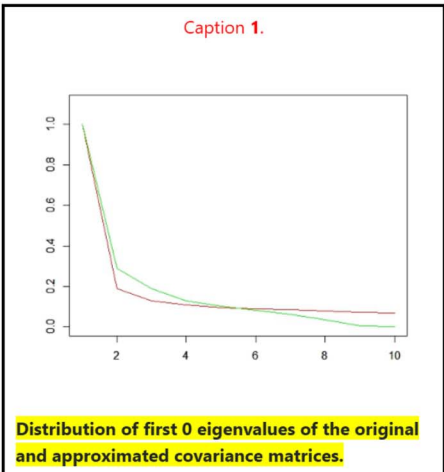


Figure 9: The drag-and-drop interface used by professional editors to rank captions for a figure [Figure source: Bulgakov (2016)].