

RADAR: A Reasoning-Guided Attribution Framework for Explainable Visual Data Analysis

Anku Rani^{1*} Aparna Garimella² Apoorv Saxena²
Balaji Vasanth Srinivasan² Paul Pu Liang¹

¹ Massachusetts Institute of Technology ² Adobe Research
ankurani@mit.edu

Abstract

Data visualizations like charts are fundamental tools for quantitative analysis and decision-making across fields, requiring accurate interpretation and mathematical reasoning. The emergence of Multimodal Large Language Models (MLLMs) offers promising capabilities for automated visual data analysis, such as processing charts, answering questions, and generating summaries. However, they provide no visibility into which parts of the visual data informed their conclusions; this black-box nature poses significant challenges to real-world trust and adoption. In this paper, we take the first major step toward evaluating and enhancing the capabilities of MLLMs to attribute their reasoning process by highlighting the specific regions in charts and graphs that justify model answers. To this end, we contribute RADAR, a semi-automatic approach to obtain a benchmark dataset comprising 1000 charts, 2000 question-answer pairs, 3599 reasoning steps, and 11,220 attribution annotations. We also introduce a method that provides attribution for chart-based mathematical reasoning. Experimental results demonstrate that our reasoning-guided approach improves attribution accuracy by up to 15 percentage points compared to baseline methods, and enhanced attribution capabilities translate to stronger answer generation, achieving high semantic similarity (BERTScore ~ 0.90) with ground truth responses. This advancement represents a significant step toward more interpretable and trustworthy chart analysis systems, enabling users to verify and understand model decisions through reasoning and attribution.

1 Introduction

Data visualizations, particularly bar charts and line charts, serve as fundamental tools for data representation and analysis across domains. These visualizations enable informed decision-making by

presenting complex numerical information in an accessible format. However, extracting insights from charts often requires sophisticated mathematical reasoning like complex trend analysis and comparative calculations (Kim et al., 2020). The emergence of Multimodal Large Language Models (MLLMs) offers promising capabilities for automated visual data analysis (Satpute et al., 2024; Srivastava et al., 2024; Ahn et al., 2024; Gupta et al., 2024), but they do so without providing any visibility into which parts of the chart informed their decision-making process (Wang et al., 2023). This poses problems for real-world trust and adoption, especially since charts and diagrams are often used in sensitive applications like business, medicine, education, and policy (Islam et al., 2024).

In this paper, we take the first major step towards evaluating and enhancing the capabilities of MLLMs to attribute their reasoning process to visual data like charts and graphs. We introduce attribution to identify and highlight key regions within charts through bounding boxes that justify final decisions, helping make the reasoning process transparent and verifiable. Previous research on attribution has largely concentrated on text-based question-answering and general visual question-answering tasks (Yue et al., 2023; Phukan et al., 2024a,b; Bohnet et al., 2022; Qi et al., 2024), which face significant limitations when applied to mathematical chart analysis. Figure 1 shows that existing methods struggle to accurately identify relevant chart regions that contribute to answers for complex mathematical questions. This gap is particularly notable as mathematical reasoning with charts requires precise identification of data points and understanding of their relationships for operations like aggregation, comparison, and trend analysis.

To bridge this gap, we introduce **RADAR**: A **Re**asoning-**Guide**D **Att**ribution Framework for Explainable Visual Data Analysis, a novel framework for attributing generated answers to specific

*A part of the work was done when the author was at Adobe Research

Question: Is the difference between blue and green graph in 2008 the maximum among years? **Answer:** No

Reasoning: In 2008, the green ("Yes") is 69%, and the blue ("No") is 26%. Difference = 43%. In 2013, the green ("Yes") is 76%, and the blue ("No") is 18%. Difference = 58%. The 2013 difference (58%) is greater than the 2008 difference (43%). The maximum difference is in 2013, not 2008.

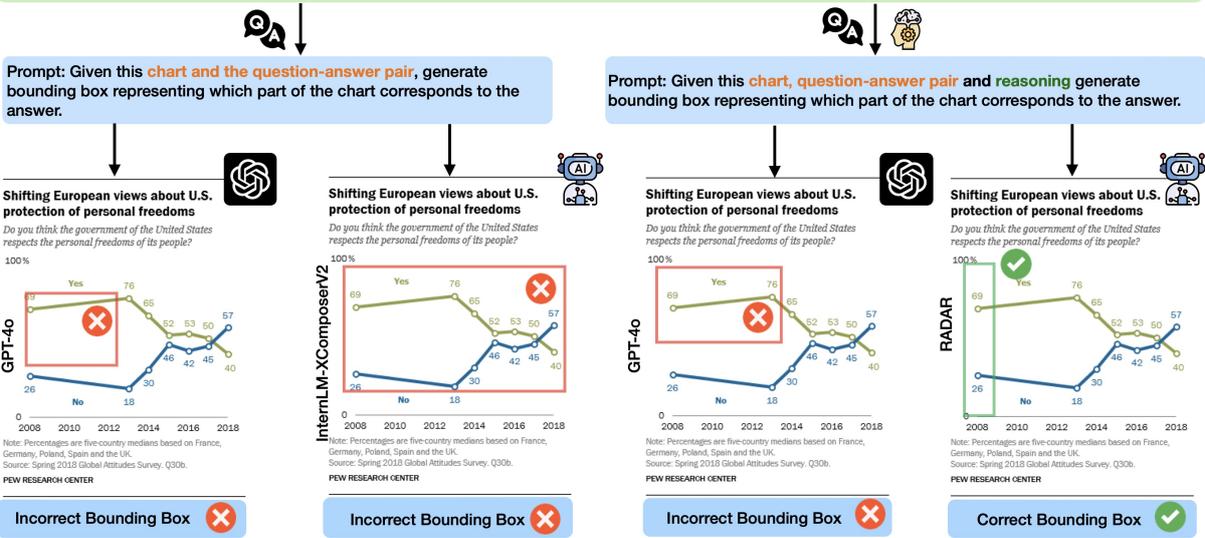


Figure 1: Comparison of attribution methods between GPT-4o & RADAR (Our method that utilizes InternLM-XComposerV2). Left: Both models receive only the chart, question, and answer as input. Right: Models additionally receive reasoning steps, leading to more precise attributions. The example shows how incorporating reasoning steps helps InternLM-XComposer2 correctly attribute the relevant data points for comparing differences between lines in 2008 and 2013, while GPT-4o struggles with accurate attribution even with reasoning provided.

regions in charts. RADAR has four main contributions:

1. We introduce the **task of attribution for mathematical question answering in charts**, addressing a critical gap in current visual mathematical question answering, as outlined in Figure 1.
2. We present a systematic **data curation strategy** that combines MLLM-generated reasoning and attribution annotations with human corrections. This results in a high-quality dataset derived from ChartQA (Masry et al., 2022), comprising annotated examples spanning line and bar chart types and mathematical operations.
3. We **propose an automatic attribution and reasoning generation method** for ChartQA and evaluate the generated attribution automatically by generating answers using feature attribution.
4. We demonstrate that high-quality **reasoning steps substantially improve attribution accuracy**, achieving up to 15 percentage point improvement over baseline methods, highlighting both the importance of reasoning in attribution tasks and identifying opportunities for further advances in reasoning generation.
5. We establish that enhanced attribution capabilities lead to more accurate answer generation, as evidenced by high BERTScore values (~ 0.90)

shown in Table 6, demonstrating the **synergistic relationship between attribution and answer quality**.

2 Related Work & Attribution Definition

Recent work has increasingly focused on attribution mechanisms for AI systems' outputs. For text systems, Bohnet et al. (2022) survey attribution methods in open-domain generation, while Phukan et al. (2024b) and Qi et al. (2024) advance answer attribution through hidden state analysis and the MIRAGE approach, respectively. Mathur et al. (2024) introduce MATSA for multi-agent table structure attribution, providing row-and-column-level attributions for claims derived from tables. For multimodal systems, Phukan et al. (2024a) extend logit lens techniques to detect visual hallucinations.

Chart-based question answering has evolved significantly, with the ChartQA dataset (Masry et al., 2022) providing a comprehensive benchmark and advances like Chart Llama (Han et al., 2023) and ChartOCR (Luo et al., 2021) improving chart understanding capabilities. Concurrent with our work, Suri et al. (2025) introduced ChartLens, which employs segmentation-based techniques (SAM and heuristics) with set-of-marks prompting for answer-

level chart attribution. While ChartLens focuses on visual grounding through segmentation, our work introduces reasoning-guided attribution with a two-level framework (answer and reasoning-level) and demonstrates that step-by-step reasoning improves attribution accuracy.

Mathematical reasoning has become crucial for chart interpretation, with Imani et al. (2023) introducing MathPrompter for multiple solution paths and Ranaldi and Freitas (2024) addressing Chain-of-Thought limitations. As Lu et al. (2023) note, mathematical reasoning serves as a key testbed for AI capabilities. While models like InternLM-XComposer2 (Dong et al., 2024) excel in multi-modal understanding, existing attribution methods face challenges with mathematical chart questions, creating a critical gap in trustworthy chart-based mathematical reasoning systems.

2.1 Attribution Definition

Chart attribution aims to identify regions of a chart that support generated answers, similar to the Grounded Visual Question Answering (VQA) approach proposed by (Phukan et al., 2024a). For mathematical chart question answering, where complex reasoning steps are essential to arrive at answers, we propose a two-level attribution framework that provides transparency not only for final answers but also for intermediate reasoning steps.

2.2 Answer-Level Attribution

Answer-Level Attribution involves visually linking chart elements to answers using bounding boxes, highlighting the specific data points that support the answer. In the leftmost chart of Figure 2, the bounding boxes highlight all data points contributing to the answer "5478". While the bounding boxes highlight all data points, the reasoning behind this calculation remains unclear without additional context. This demonstrates why incorporating reasoning steps becomes crucial for questions involving mathematical operations, where the path to the answer is as important as the answer itself.

2.3 Reasoning-Level Attribution

Reasoning-Level Attribution involves visually linking each reasoning step to relevant chart regions, creating a traceable connection between the reasoning process and visual elements. As shown in Figure 2, the 2nd, 3rd and 4th chart represents each of the reasoning steps. This granular-level attribution approach enhances trust in the system by

Model	Answer is Correct (Human Annotated)	Reasoning is Correct (Human Annotated)
Gpt-4o	58%	49%
Gpt-4v	64%	45%
Claude-3.5-sonnet	96%	75%

Table 1: Benchmarking performance based on Visual Question Answering and Visual Question Reasoning. **Green** reflects the highest score, **Orange** reflects the mid ranged score and **Red** reflects the lowest score.

making both final answers and the reasoning process transparent and verifiable against the source chart.

Additional examples for both answer level and reasoning level attribution are present in the Appendix section A.

3 Dataset Curation

Currently, no datasets exist that provide reasoning steps for chart question answering or attribution annotations for mathematical chart QA. To address this gap, we first examine existing model capabilities before developing a semi-automatic annotation strategy.

3.1 Reasoning Capabilities of MLLMs for Charts

Based on performance evaluations in the Polymath benchmark (Gupta et al., 2024), we select Claude 3.5 Sonnet and GPT-4o as our primary models for analysis. We also include GPT-4v for its vision capabilities. To assess the performance on reasoning generation, we randomly select 100 examples from the ChartQA dataset (Masry et al., 2022), each containing a chart, question, and answer triple. Charts and questions are passed to these models as input, and they are prompted to generate answers and reasoning. These answers and reasoning are annotated by human annotators on whether they are correct or not, and the results are presented in Table 1.

As shown in Table 1, while Claude 3.5 Sonnet demonstrates strong answer generation (96% accuracy), its reasoning capabilities show significant room for improvement (75% accuracy). Other models perform notably worse, with reasoning accuracies below 50%.

A more detailed error analysis, including word clouds of failure patterns in appendix section B figure 12, taxonomy in Appendix section B Figure 11 and specific examples is provided in the appendix section B Figure 13. Notably, providing

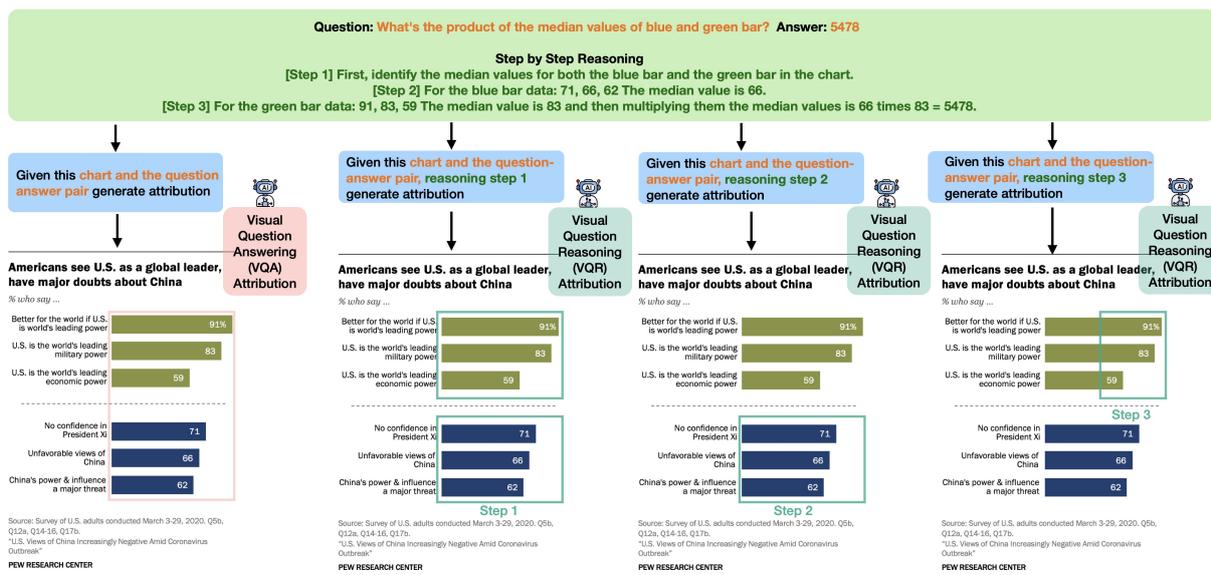


Figure 2: RADAR’s attribution process for mathematical reasoning on charts. Left to right: (1) VQA attribution highlights all relevant data bars needed for computing the final answer, (2) Step 1 attribution identifies the specific bars needed to find median values, (3) Step 2 attribution focuses on the blue bars for calculating their median (66), and (4) Step 3 attribution shows the green bars used for the final multiplication step ($66 \times 83 = 5478$). The attribution represents how our framework traces both the final answer & intermediate reasoning steps.

correct answers alongside questions reduced reasoning failures from 51% to 25%, suggesting the potential for improved performance through better model guidance.

3.2 Task Setup

Our goal is to obtain three types of annotations for chart-based mathematical attribution: (1) reasoning steps for given chart-question-answer triples, (2) answer attribution and (3) reasoning step attribution. Rather than annotating from scratch, we developed a semi-automatic approach leveraging Claude 3.5 Sonnet’s capabilities to generate initial annotations for human correction.

We recruited two qualified annotators through the Upwork platform¹ after an initial screening of three candidates using 100 sample data points. The entire annotation process went for 120 hours and each annotator was paid 15 USD hourly.

For attribution annotation, we employed the VGG Image Annotator platform², which provides an intuitive interface for drawing bounding boxes and mapping them to textual reasoning steps. Screenshots from the annotation interface, more details on initial screening and examples of such annotations are provided in the appendix section A Figure 5 and 9.

¹<https://www.upwork.com>

²<https://annotate.officialstatistics.org/>

Stage 1: Reasoning Validation and Correction. In Stage 1, annotators perform reasoning validation through three key steps: (1) correction by reviewing Claude-3.5-sonnet generated reasoning for chart-question-answer triples, (2) providing a binary correctness assessment (Yes/No) for each triple, and (3) categorizing errors in incorrect reasoning (such as color mismatches or illogical conclusions) while supplying corrected reasoning when the original is found to be inaccurate.

Stage 2: Answer Attribution. For each chart-question-answer triple, annotators draw bounding boxes using the VGG image annotator indicating chart regions supporting the answer. Figure 5 demonstrates this process, showing both input and resulting annotations.

Stage 3: Reasoning Attribution. Using the validated reasoning from Stage 1, annotators are instructed to provide bounding boxes using the VGG image annotator for each reasoning step. As shown in Figure 9, each statement (e.g., "orange line represents unfavorable") is linked to relevant chart regions.

3.3 Data Annotation & Analysis

To ensure annotation quality, we conducted initial screening to select mathematically proficient annotators, measured inter-annotator agreement using Kappa score (Cohen, 1960), and had authors man-

Chart Type	Stage 1 [Kappa Score]	Stage 2 [IOU Score]	Stage 3 [IOU Score]
Line	0.82	0.52	0.56
Bar	0.92	0.57	0.64

Table 2: Across three annotation stages, we observed high agreement for reasoning validation (Kappa > 0.8), moderate agreement for answer-based attribution (IOU 0.5), and improved agreement with reasoning-based attribution (IOU 0.56-0.64), indicating reasoning steps enhance annotation consistency. **Green** reflects the highest score, **Orange** reflects the mid ranged score and **Red** reflects the lowest score.

Chart Type	Charts	QA Pairs	Reasoning Steps	Attributed Regions (QA)	Attributed Regions (Reasoning)
Line	500	1000	1773	1465	2691
Bar	500	1000	1826	2627	4437
Total	1000	2000	3599	4092	7128

Table 3: RADAR Dataset contains 1,000 charts (500 line, 500 bar) with 2,000 QA pairs, 3,599 reasoning steps, and 11,220 total bounding box annotations (4,092 answer-level + 7,128 reasoning-level).

ually verify a sample of annotations. This semi-automatic approach significantly reduced annotation effort while maintaining high quality through human validation and correction. More details on the Inter-annotator Agreement consisting of Kappa score and Intersection Over Union (IOU) score (Rezatofighi et al., 2019) calculation formula are present in the Appendix section B.1.

3.4 Data Analysis

After the annotation process, the key statistics about the data are summarized in the table 3. Table 3 shows the breakdown of the dataset by chart type. There are a total of 1000 charts, consisting of 500 line charts and 500 bar charts. For each chart, there are 2 QA pairs, resulting in a total of 2000 QA pairs. Additionally, the annotators identified a total of 3599 reasoning steps across all the charts (stage 1). The table also shows the number of image regions that were attributed to the QA-based annotations (stage 2) and the reasoning-based annotations (stage 3). There are a total of 1,000 charts with 2,000 QA pairs, 3,599 reasoning steps, and 11,220 bounding box annotations. For line charts, there are 1465 QA-based and 2691 reasoning-based attributed regions, while for bar charts, there are 2627 QA-based and 4437 reasoning-based attributed regions.

4 RADAR: Proposed Method

Our proposed method, RADAR, addresses the challenge of attributing mathematical reasoning in charts through a two-stage pipeline. Given a chart-question-answer triple, we first generate step-by-step reasoning using InternLM-XComposer2 model (figure 3), then leverage these reasoning steps along with the chart, question, and answer to produce attribution bounding boxes for both the final answer and intermediate reasoning steps. For reasoning generation, we utilize InternLM-XComposer2³ output. The model architecture incorporates a vision encoder (CLIP ViT-Large (Radford et al., 2021)) that processes charts into a 35×35 grid (1225 visual tokens) and maps them to a shared 4096-dimensional embedding space with text from InternLM-2 (Cai et al., 2024). Using (Dong et al., 2024), we employ Partial LoRA (Hu et al., 2021) (PLoRA), which applies additional trainable parameters specifically to visual tokens while preserving the base 7B-parameter language model’s capabilities.

We utilize (Phukan et al., 2024a)’s findings on attribution, and extract hidden states from layer 16, which empirically provides optimal semantic representations for our task. The attribution mechanism employs a GPU-accelerated sliding window approach, efficiently processing window configurations from 3×3 to 35×35 patches through normalized patch embedding averaging and cosine similarity metrics between textual descriptions and visual regions.

5 Experiments

We conduct experiments on the curated dataset presented in section 3.4. We experiment on this dataset for three tasks i.e. (i) Attribution based on Visual Question Answering (VQA), (ii) Attribution based on Visual Question Reasoning (VQR), and (iii) Answer generation based on feature attribution.

5.1 Baselines

We evaluate RADAR against three state-of-the-art MLLMs: GPT-4o, GPT-4v, and Claude 3.5 Sonnet. For each baseline, we test both zero-shot and few-shot prompting strategies for two tasks: attribution based on answer attribution and reasoning attribution. We evaluate the generated attribution using answer generation using feature attribution.

³<https://github.com/InternLM/InternLM-XComposer>

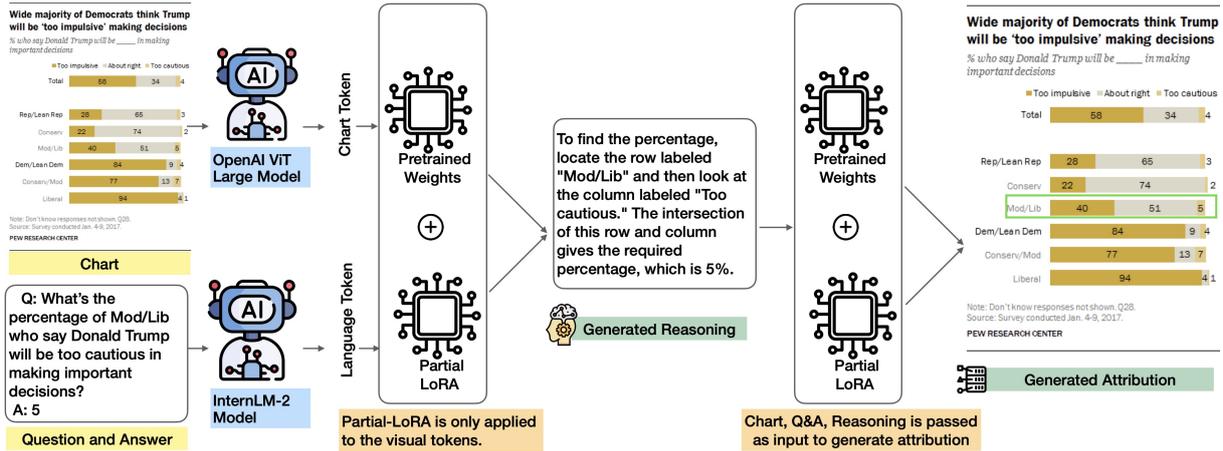


Figure 3: Overview of our proposed RADAR framework for reasoning and attribution generation using InternLM-XComposer2. Given a chart and Q&A pair, the model processes visual tokens through CLIP ViT-Large and applies Partial-LoRA for chart-specific feature adaptation, while InternLM-2 processes textual inputs. The final output provides both answers and reasoning. This process is repeated by passing Chart, Q&A pair, and generated reasoning to InternLM-XComposer2 to generate attribution.

Answer Attribution (VQA): Models must identify relevant chart regions using bounding boxes that support their answers to specific questions.

Reasoning Attribution (VQR): Models must attribute their mathematical reasoning steps using bounding boxes to specific chart elements. Each reasoning step has a different granular attribution as described in VQR steps 1, 2, and 3 of fig 2. We collect attribution results through API calls to GPT-4o⁴, GPT-4v⁵, and Claude 3.5 Sonnet⁶. Since these models cannot directly output chart images with bounding boxes, we design prompts to obtain coordinates of the bounding boxes in the format of X1, Y1, X2, and Y2.

For **Answer Attribution (VQA)**, we explored two prompting strategies for answer attribution: (1) Zero-shot, where we used a base prompt without examples, and (2) Few-shot, where we included example cases in the prompt. Both prompting approaches are detailed in Appendix section C.2.

For **Reasoning Attribution (VQR)**, we conducted both zero-shot and few-shot prompting for our visual question reasoning attribution. The detailed prompts are available in Appendix section C, with zero-shot prompts shown in figure 16 and few-shot prompts in figure 17. Input charts were base64 encoded for prompt delivery and decoded after collecting bounding box coordinates for subsequent analysis.

⁴GPT-4o: "GPT-4o", "2023-05-15"

⁵GPT-4v: "gpt-4-vision-preview", "2023-07-01-preview"

⁶Claude 3.5 Sonnet: "claude-3-5-sonnet-20240620-v1:0"

5.2 RADAR-Automatic Reasoning Step Generation

For automatic reasoning step generation, we use the pre-trained InternLM-XComposer2 model. We leverage the Partial LoRA framework from (Dong et al., 2024) due to its effectiveness in preserving language capabilities while adapting to visual inputs. Given a chart-question-answer triple (C, Q, A) , our goal is to generate reasoning steps R that explain the answer derivation while maintaining alignment with visual elements.

Following the existing Partial LoRA architecture, we process inputs $x = [x_v, x_t]$, where x_v represents visual tokens from the chart processed through CLIP ViT-Large, and x_t represents the concatenated question-answer tokens. The output features are computed as follows:

$$\hat{x} = [\hat{x}_v, \hat{x}_t] \quad (1)$$

where \hat{x}_t follows the standard language model path, and \hat{x}_v incorporates visual adaptation through the Partial LoRA matrices.

This approach enables our model to generate step-by-step reasoning by utilizing chart-specific visual features while maintaining strong language capabilities, producing coherent explanations that explicitly reference chart elements, and describing the mathematical operations needed to arrive at the answer. The generated reasoning provides a transparent explanation of the answer derivation process, which is then used to guide our attribution mechanism for identifying relevant chart regions.

5.3 RADAR-Answer Generation Using Feature Attribution

To automatically evaluate RADAR without involving humans, we implemented a feature attribution mechanism that emphasizes the salient regions of the chart. Specifically, we identify the bounding box containing the ground truth answer and mask the remainder of the image by setting all pixels outside this region to zero. This preprocessed chart, along with the original question, is then passed to RADAR for answer extraction as described in Appendix section C.4 Figure 20. We evaluate the effectiveness of this approach by computing a similarity score between the model’s extracted answer and the ground truth which leads to automatic evaluation of RADAR.

5.4 Metrics

Reasoning Generation Task

In figure 3, automatic reasoning is generated by our proposed method RADAR. We also have human-annotated reasoning collected from stage 1 of Task setup as described in section 3.2. We evaluate the generated reasoning using BERTScore (Zhang et al., 2019) & Semantic Textual Similarity(STS) (Agirre et al., 2012) report the scores in Table 4. Additional details are present in Appendix section C.4.

Attribution Generation Task

We use **Multiple Box IOU Score** by extending the IOU score (Rezatofighi et al., 2019) from single to multiple bounding boxes representing the utilization of pixels. While traditional single-box IOU simply computes the overlap between two rectangular regions, multiple-box IOU accounts for complex spatial relationships between sets of boxes, including internal overlaps within each set. Our implementation solves this challenge by converting box sets into binary masks, converting a geometric problem into a pixel-wise operation. This mask-based approach automatically handles cases where boxes within a set overlap, eliminating the need for complex geometric calculations of polygon intersections. The result is a Multi-Box IOU metric that maintains the intuitive interpretation of IOU while handling multiple boxes, helping to predict multiple regions for a single object or class.

$$IOU_{multi} = \frac{|\mathcal{M}_p \cap \mathcal{M}_{gt}|}{|\mathcal{M}_p \cup \mathcal{M}_{gt}|} \quad (2)$$

where \mathcal{M}_p and \mathcal{M}_{gt} are binary masks generated from the sets of predicted and ground truth

Chart Type	BERTScore			Semantic Textual Similarity
	Avg Precision	Avg Recall	Avg F1	
Line	0.892	0.885	0.889	0.738
Bar	0.892	0.886	0.889	0.742

Table 4: Comparison of reasoning quality between RADAR-generated and human-annotated explanations. The evaluation uses two metrics: BERTScore and Semantic Textual Similarity (STS) converted to a percentage. Both chart types show consistent performance with BERTScore values around 0.892 and STS scores above 0.738, indicating strong alignment between machine-generated and human reasoning. **Green** reflects the highest and **Orange** reflects the lowest.

bounding boxes respectively.

Answer Generation Task

Attribution generated by RADAR as described in Figure 4 is passed through the method described in section 5.3 to generate answers. We evaluate the generated answer using BERTScore (Zhang et al., 2019) & Semantic Textual Similarity (STS) (Agirre et al., 2012) with human-annotated ones and report the scores in Table 6.

6 Results and Discussion

Reasoning Generation Task

We report average BERTScore & STS in table 4 between RADAR-generated reasoning and human-annotated reasoning. Results show strong alignment across chart types (precision: 0.8928, recall/F1: ~ 0.89) with comparable STS scores (~ 0.74), demonstrating reliable reasoning capabilities that mirror human annotations.

Attribution Generation Task

Table 5 presents the Multi Box IOU scores across different models and prompting strategies. RADAR demonstrates substantial improvements: 13.1% (2.6 \rightarrow 15.7) and 12.5% (2.8 \rightarrow 15.3) for VQA tasks in line and bar charts respectively, while showing 8.5% (12.2 \rightarrow 3.7) & 4.3% (3.9 \rightarrow 8.2) improvement for VQR tasks with automated reasoning. The improvements increase further with human-validated reasoning, particularly for VQR tasks (11.1% for line charts, 17.6% for bar charts). While these results demonstrate significant progress, the best-performing variant achieves IOU scores % of 15% - 20%, indicating room for improvement through better attribution systems and MLLMs.

Answer Generation Task

We take average STS and BERTScore to com-

Model	VQA IOU_{multi}		VQR IOU_{multi}	
	Line	Bar	Line	Bar
GPT-4o (zero-shot)	2.6%	2.8%	2.5%	2.1%
GPT-4o (few-shot)	2.0%	2.2%	2.2%	1.9%
GPT-4v (zero-shot)	1.6%	1.9%	2.1%	2.3%
GPT-4v (few-shot)	1.4%	1.7%	2.2%	2.4%
Claude 3.5 (zero-shot)	2.4%	2.5%	3.2%	3.5%
Claude 3.5 (few-shot)	2.5%	2.1%	3.7%	3.9%
RADAR (Automated Reasoning)	15.7%	15.3%	12.2%	8.2%
RADAR (Human Reasoning)	15.7%	15.3%	13.6%	19.7%

Table 5: Attribution performance comparison across different models and settings. Performance is measured using IOU_{multi} scores for both VQA and VQR tasks. The scores are converted to a percentage. Green reflects the highest, Orange reflects the mid-range and Red reflects the lowest.

Chart Type	Bertscore		Semantic Textual Similarity	
	Avg Precision	Avg Recall	Avg F1	
Line	89.3%	84.3%	86.6%	50.5%
Bar	90.3%	86.1%	88.1%	51.9%

Table 6: Comparison of textual similarity between RADAR-generated answers using feature attribution and ground truth using Bertscore and STS metrics across line and bar charts, demonstrating consistent performance across visualization formats. The scores are converted to a percentage. Green reflects the highest and Orange reflects the lowest.

pare generated answers using feature attribution and the ground truth. Table 6 shows the effectiveness of feature attribution in RADAR for answer generation across line and bar charts. Bar charts consistently outperform line charts across all metrics, with higher Bertscore values in precision (0.9036), recall (0.8612), and F1 (0.8813). The moderate STS scores (~ 0.51) and marginal performance differences between chart types suggest RADAR’s feature attribution approach is robust across visualization formats.

7 Generalization & Scalability

We extracted 1068 pie charts with their Question, Answer pair from the ChartQA dataset and processed them through three steps: (i) reasoning generation, (ii) attribution generation, and (iii) answer generation using feature attribution as described in Figure 4. The generated answers, when evaluated to test the attribution system’s accuracy, achieve an average BERTScore of ~ 0.9 and average STS of ~ 0.5 . This fully automated approach demonstrates scalability and generalizes effectively across various chart types when provided with charts and

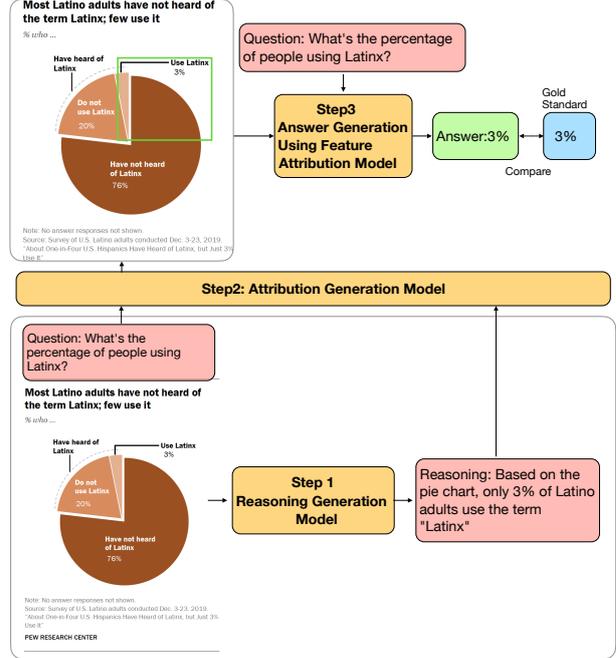


Figure 4: Extension of our proposed framework to pie charts, illustrating the framework’s adaptability across different visualization formats.

questions.

Beyond the evaluated chart types, our framework shows promise for extension to other visualization formats. From a computational scalability perspective, the framework processes charts efficiently at inference time, requiring approximately 1-2 seconds per chart-question pair on standard GPU hardware. This efficiency positions RADAR as a practical solution for large-scale chart understanding applications, with the framework maintaining consistent performance across 1,000+ diverse charts without task-specific fine-tuning.

8 Conclusion

In this paper, we presented a novel framework for chart attribution that combines visual and mathematical reasoning capabilities. Our primary contributions include (i) the formalization of chart attribution for mathematical question answering and reasoning tasks, (ii) a systematic data curation strategy that combines MLLM-generated reasoning with human corrections for reliable attribution annotation, and (iii) a framework that utilizes automatic reasoning steps to improve attribution accuracy. While our approach demonstrates significant improvements over baselines, opportunities remain for enhancing reasoning generation, extending support for complex chart types, and integrating with

downstream applications. Our framework provides a foundation for building more trustworthy and interpretable AI systems for mathematical reasoning tasks, paving the way for chart-based systems that can better explain their decision-making processes.

9 Limitations

While our framework demonstrates promising results for chart attribution through automated reasoning, several important limitations and areas for discussion emerge from our study:

Restrictive Prompting: To produce bounding boxes, we use a highly restrictive prompting approach. We instruct the model to generate bounding boxes as a list of coordinate tuples in the format $(X1, Y1, X2, Y2)$. However, research works on restrictive prompting (Tam et al., 2024) has found that using overly restrictive prompts can lead to notable decreases in model performance compared to less constrained prompting techniques. Therefore

Attribution Task is Challenging for Humans: As shown in Table 2, the attribution task proved difficult for human annotators. In stages 2 and 3 of the annotation process, the agreement percentages ranged from just 52% to 64%. These relatively low levels of agreement underscore the inherent challenge of the attribution task, even for human raters with domain expertise.

Reasoning Quality Dependencies: Our attribution system’s performance depends on the quality of generated reasoning steps. While fine-tuning InternLM-XComposer2 may improve reasoning generation, complex mathematical operations, and multi-step calculations still present challenges, potentially affecting attribution accuracy. We discuss failure cases for reasoning in fig 12.

Computational Requirements: The sliding window mechanism used for attribution, while effective, requires significant computational resources, especially for high-resolution charts or when processing multiple reasoning steps. This may impact the system’s practicality in real-time applications.

Human Validation Process: While our data curation strategy employs human validation to ensure quality, the subjectivity in reasoning annotation and attribution marking can introduce inconsistencies. The inter-annotator agreement scores suggest room for improvement in standardizing the validation process.

Model Architecture Limitations: The current approach relies on layer 16 hidden states of InternLM-XComposer2, which may not capture all relevant features for attribution. Alternative architectural choices or multi-layer approaches could potentially yield better results.

These limitations point to several promising di-

rections for future research, including more robust reasoning generation mechanisms, efficient attribution algorithms, and improved validation methodologies.

10 Ethics Statement

We acknowledge several ethical considerations in our development of chart attribution systems. First, we prioritized transparency by openly documenting our methodology, model limitations, and potential biases in both reasoning generation and attribution accuracy. All training data was properly sourced from public datasets with appropriate licensing, and our human annotation process followed fair labor practices, including equitable compensation (\$15/hour) and clear guidelines. While our system aims to improve accessibility and understanding of quantitative information through transparent reasoning steps, we recognize potential risks of misuse, such as the automated generation of misleading chart interpretations. We recommend deploying this technology with appropriate human oversight in high-stakes scenarios and maintaining regular audits for systematic biases. Our goal is to advance chart interpretation capabilities while implementing safeguards that protect against potential misuse and ensure the technology serves its intended purpose of making quantitative information more accessible and understandable to diverse user groups.

References

- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. [SemEval-2012 task 6: A pilot on semantic textual similarity](#). In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393, Montréal, Canada. Association for Computational Linguistics.
- Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. 2024. Large language models for mathematical reasoning: Progresses and challenges. *arXiv preprint arXiv:2402.00157*.
- Bernd Bohnet, Vinh Q Tran, Pat Verga, Roe Aharoni, Daniel Andor, Livio Baldini Soares, Massimiliano Ciaramita, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, et al. 2022. Attributed question answering: Evaluation and modeling for attributed large language models. *arXiv preprint arXiv:2212.08037*.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaxing Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yining Li, Hongwei Liu, Jiangning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingtong Xiong, Chao Xu, Ruiliang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zhang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. 2024. [Internlm2 technical report](#).
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Bin Wang, Linke Ouyang, Xilin Wei, Songyang Zhang, Haodong Duan, Maosong Cao, et al. 2024. [Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model](#). *arXiv preprint arXiv:2401.16420*.
- Himanshu Gupta, Shreyas Verma, Ujjwala Anantheswaran, Kevin Scaria, Mihir Parmar, Swaroop Mishra, and Chitta Baral. 2024. [Polymath: A challenging multi-modal mathematical reasoning benchmark](#).
- Yucheng Han, Chi Zhang, Xin Chen, Xu Yang, Zhibin Wang, Gang Yu, Bin Fu, and Hanwang Zhang. 2023. [Chartllama: A multimodal llm for chart understanding and generation](#). *arXiv preprint arXiv:2311.16483*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *arXiv preprint arXiv:2106.09685*.
- Shima Imani, Liang Du, and Harsh Shrivastava. 2023. [Mathprompter: Mathematical reasoning using large language models](#). *arXiv preprint arXiv:2303.05398*.
- Mohammed Saidul Islam, Raian Rahman, Ahmed Masry, Md Tahmid Rahman Laskar, Mir Tafseer Nayeem, and Enamul Hoque. 2024. [Are large vision](#)

- language models up to the challenge of chart comprehension and reasoning? an extensive investigation into the capabilities and limitations of vlms. *arXiv preprint arXiv:2406.00257*.
- Dae Hyun Kim, Enamul Hoque, and Maneesh Agrawala. 2020. [Answering questions about charts and generating visual explanations](#). In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20*, page 1–13, New York, NY, USA. Association for Computing Machinery.
- Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and Kai-Wei Chang. 2023. [A survey of deep learning for mathematical reasoning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14605–14631, Toronto, Canada. Association for Computational Linguistics.
- Junyu Luo, Zekun Li, Jinpeng Wang, and Chin-Yew Lin. 2021. [Chartocr: Data extraction from charts images via a deep hybrid framework](#). In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1917–1925.
- Ahmed Masry, Do Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. [ChartQA: A benchmark for question answering about charts with visual and logical reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279, Dublin, Ireland. Association for Computational Linguistics.
- Puneet Mathur, Alexa Siu, Nedim Lipka, and Tong Sun. 2024. [MATSA: Multi-agent table structure attribution](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 250–258, Miami, Florida, USA. Association for Computational Linguistics.
- Anirudh Phukan, Divyansh, Harshit Kumar Morj, Vaishnavi, Apoorv Saxena, and Koustava Goswami. 2024a. [Beyond logit lens: Contextual embeddings for robust hallucination detection & grounding in vlms](#).
- Anirudh Phukan, Shwetha Somasundaram, Apoorv Saxena, Koustava Goswami, and Balaji Vasani. 2024b. [Peering into the mind of language models: An approach for attribution in contextual question answering](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 11481–11495, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Jirui Qi, Gabriele Sarti, Raquel Fernández, and Arianna Bisazza. 2024. [Model internals-based answer attribution for trustworthy retrieval-augmented generation](#). *arXiv preprint arXiv:2406.13663*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *International Conference on Machine Learning*.
- Leonardo Ranaldi and Andre Freitas. 2024. [Aligning large and small language models via chain-of-thought reasoning](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1812–1827, St. Julian’s, Malta. Association for Computational Linguistics.
- Hamid Reza Tofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. 2019. [Generalized intersection over union: A metric and a loss for bounding box regression](#). In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666.
- Ankit Satpute, Noah Gießing, André Greiner-Petter, Moritz Schubotz, Olaf Teschke, Akiko Aizawa, and Bela Gipp. 2024. [Can llms master math? investigating large language models on math stack exchange](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2316–2320.
- Pragya Srivastava, Manuj Malik, Vivek Gupta, Tanuja Ganu, and Dan Roth. 2024. [Evaluating llms’ mathematical reasoning in financial document question answering](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3853–3878.
- Manan Suri, Puneet Mathur, Nedim Lipka, Franck Dernoncourt, Ryan A. Rossi, and Dinesh Manocha. 2025. [ChartLens: Fine-grained visual attribution in charts](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22447–22462, Vienna, Austria. Association for Computational Linguistics.
- Zhi Rui Tam, Cheng-Kuang Wu, Yi-Lin Tsai, Chieh-Yen Lin, Hung yi Lee, and Yun-Nung Chen. 2024. [Let me speak freely? a study on the impact of format restrictions on performance of large language models](#).
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. 2023. [Scientific discovery in the age of artificial intelligence](#). *Nature*, 620(7972):47–60.
- Xiang Yue, Boshi Wang, Ziru Chen, Kai Zhang, Yu Su, and Huan Sun. 2023. [Automatic evaluation of attribution by large language models](#). *arXiv preprint arXiv:2305.06311*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with bert](#). *ArXiv, abs/1904.09675*.

Appendix

This section provides additional examples to assist in the understanding and interpretation of the research work presented.

Section A: Attribution Definition and Examples

Section B: Data Sources and Compilation

Section C: Experiments

A Attribution Definition & Examples

Expanding upon the definition of attribution described in section 2.1, we present additional examples for both Visual Question Answering-based attribution and Visual Question Reasoning-based attribution.

A.1 Examples of VQA based Attribution

Examples of attribution based on visual question answering are present in fig 5 and 6.

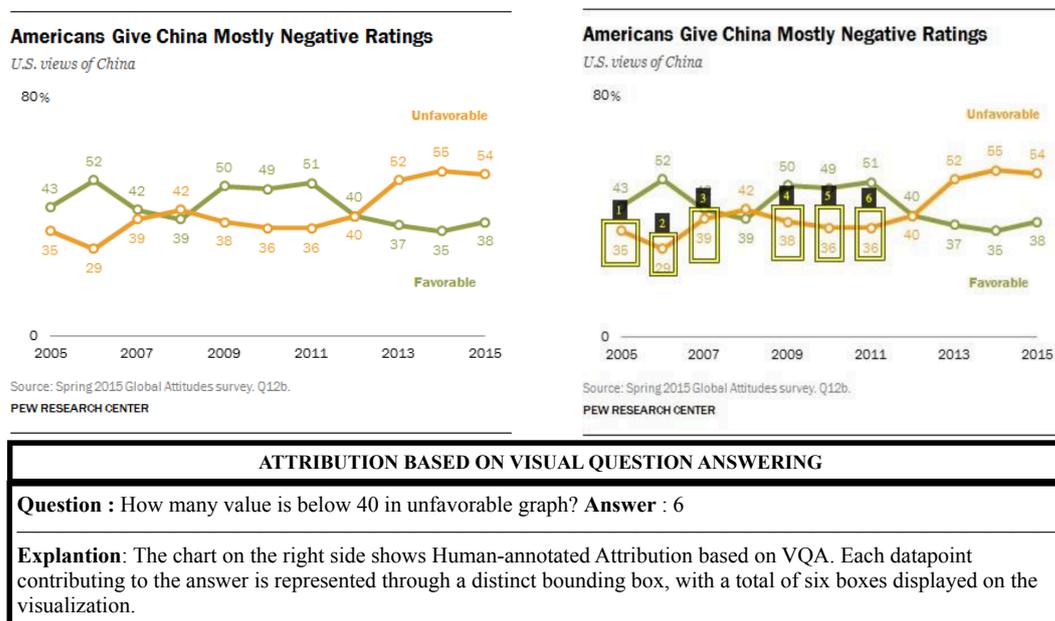


Figure 5: Example of VQA-based attribution from our dataset. Given the question "How many values are below 40 in the Unfavorable graph?" and answer "6", the annotation shows six bounding boxes identifying data points on the "Unfavorable" line that fall below 40%. Each bounding box highlights a specific chart element that directly supports the answer, demonstrating answer-level attribution.

A.2 Examples of VQR based Attribution

Examples of attribution based on visual question reasoning are present in fig 7, 8, and fig 9. The following figures also reflect the snapshot from the annotation platform.

B Data Sources and Compilation

B.1 Data Annotation

To ensure annotation quality, we conducted initial screening to select mathematically proficient annotators, measured inter-annotator agreement using Kappa score (Cohen, 1960), and had authors manually verify a sample of annotations. This semi-automatic approach significantly reduced annotation effort while maintaining high quality through human validation and correction.

The agreement score for stage 1 is based on whether the reasoning is correct and is represented by Kappa score (Cohen, 1960). Kappa score is defined as a measure of inter-rater agreement for categorical

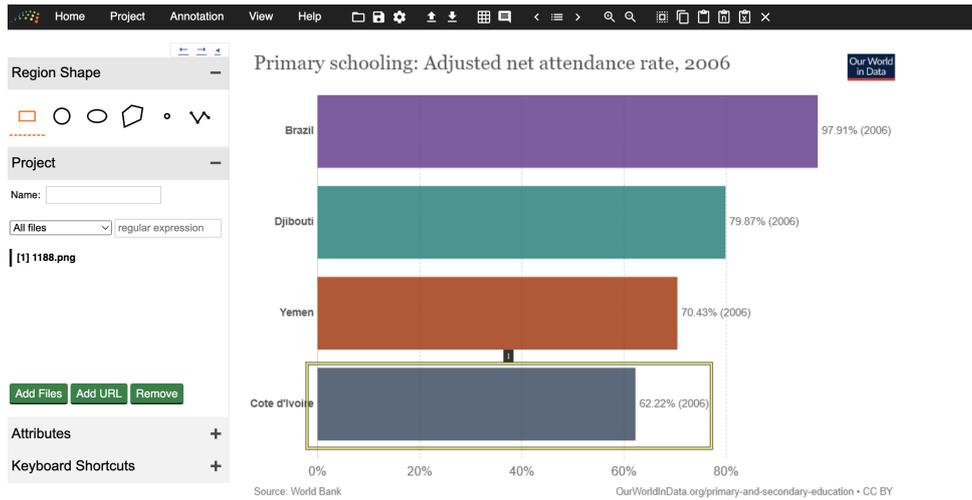


Figure 6: Screenshot from the attribution annotation platform showing the question-answering workflow. Given a chart, question ("What's the value of the smallest bar?"), and answer ("62.22%"), annotators draw bounding boxes highlighting the relevant chart regions. These bounding boxes constitute the answer-level attribution annotations used in our dataset.

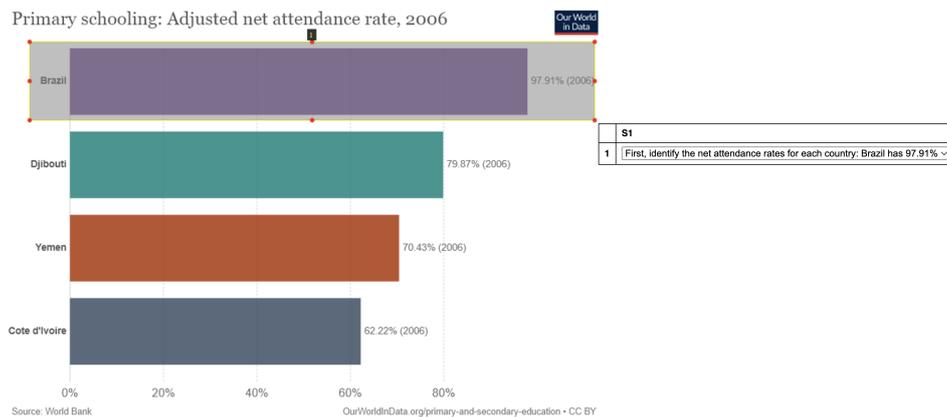


Figure 7: This figure shows how sentence-level reasoning is attributed in our dataset. Annotators are provided with a chart, question = What's the average of Yemen and Brazil?, answer= 84.17, and reasoning = "First, identify the net attendance rates for each country: Brazil has 97.91%. Yemen has 70.43%. Next, sum these values: 97.91 + 70.43 = 168.34. Then, divide by the number of countries to find the average: 168.34 / 2 = 84.17". The first reasoning statement "First, identify the net attendance rates for each country: Brazil has 97.91%." is directly linked to corresponding chart elements, ensuring each step of the mathematical reasoning process is grounded in the chart's components.

items, taking into account the agreement occurring by chance. The Kappa score is defined mathematically as:

$$\kappa = \frac{p_o - p_e}{1 - p_e}$$

where:

p_o is the observed agreement between the two annotators p_e is the expected agreement by chance
The observed agreement p_o is calculated as:

$$p_o = \frac{a + d}{a + b + c + d}$$

where: a is the number of cases where both annotators agreed on "yes" b is the number of cases where

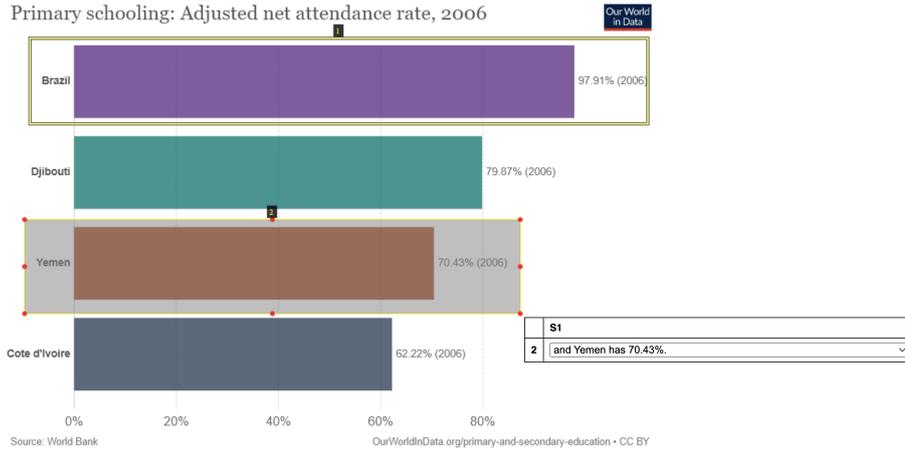
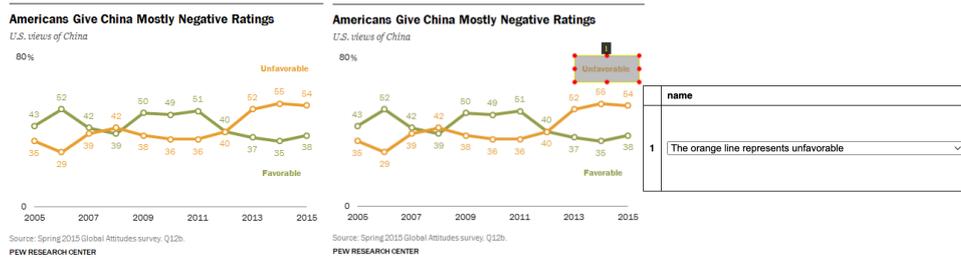


Figure 8: This figure shows how sentence-level reasoning is attributed in our dataset. Annotators are provided with a chart, question = What's the average of Yemen and Brazil?, answer= 84.17, and reasoning = "First, identify the net attendance rates for each country: Brazil has 97.91%. Yemen has 70.43%. Next, sum these values: 97.91 + 70.43 = 168.34. Then, divide by the number of countries to find the average: 168.34 / 2 = 84.17". The second reasoning statement "Yemen has 70.43%" is directly linked to corresponding chart elements, ensuring each step of the mathematical reasoning process is grounded in the chart's components.



ATTRIBUTION BASED ON VISUAL QUESTION REASONING	
Question: What is the value of unfavorable in the year 2015? Answer : 54	
Reasoning generated by Claude 3.5 Sonnet followed by correct annotation by human: The orange line represents unfavorable. For value that corresponds to the year 2015 in orange line is 54.	
Explanation: The chart on the right side shows human annotation based on reasoning steps. Given each reasoning step, we attribute it on a sentence level. For example, the first sentence i.e. orange line represents unfavorable is attributed as displayed in the chart.	

Figure 9: This figure shows how step-wise reasoning is attributed in our dataset. The first reasoning statement, "The orange line represents unfavorable," is directly linked to corresponding chart elements, ensuring each step of the mathematical reasoning process is grounded in the chart's components.

the first annotator said "yes" and the second said "no" c is the number of cases where the first annotator said "no" and the second said "yes" d is the number of cases where both annotators agreed on "no"

The expected agreement p_e is calculated as:

$$p_e = \frac{(a + b)(a + c) + (c + d)(b + d)}{(a + b + c + d)^2}$$

For stage 2 and stage 3, the Intersection over Union (IOU) score (Rezatofighi et al., 2019) was calculated. IOU score is a measure of the overlap between the bounding box drawn by annotator 1 and bounding box drawn by annotator 2, defined as the ratio of the area of intersection to the area of union. The IOU score is mathematically defined as:

$$IOU = \frac{\text{Area of Intersection}}{\text{Area of Union}}$$

5670

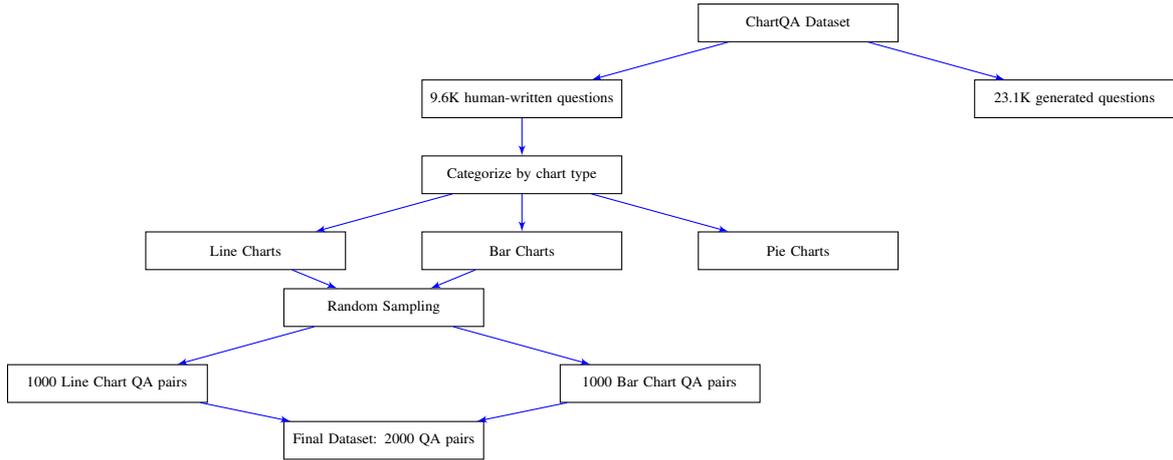


Figure 10: Data Compilation Process Flowchart. From the ChartQA dataset containing both human-written (9.6K) and generated (23.1K) questions, we focus on the human-written subset for quality assurance. We categorize these by chart type (line, bar, and pie charts), then use random sampling to create a balanced final dataset of 2000 QA pairs, comprising 1000 pairs each for line and bar charts.

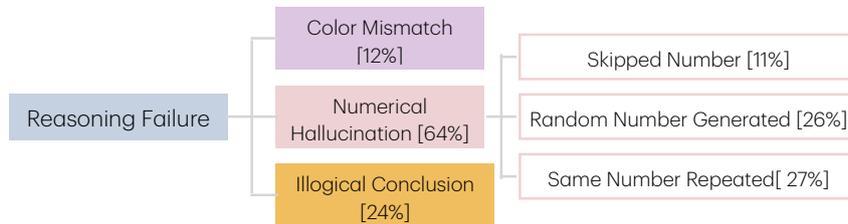


Figure 11: Taxonomy of Failure Cases that represents the categories of reasoning failure.

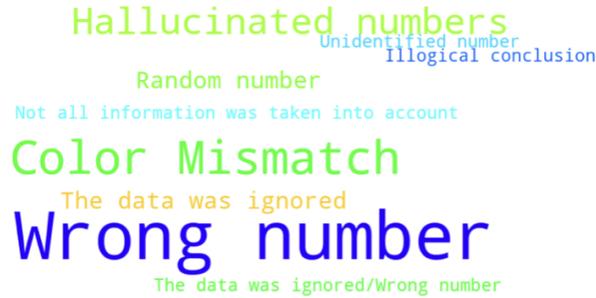


Figure 12: Word cloud representing the annotated labels for reasoning failure. These annotated labels are hallucinated numbers, Illogical conclusions, Color mismatch, data points ignored, etc.

where the area of intersection is the overlapping area between the predicted bounding box and the ground truth bounding box, and the area of union is the total area covered by both bounding boxes.

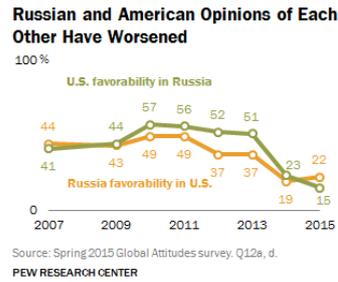
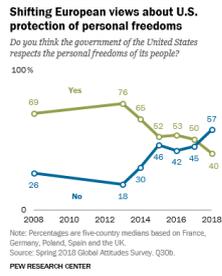
Let's denote the predicted bounding box as B_p and the ground truth bounding box as B_g . Then, the IOU score can be calculated as:

$$\text{IOU} = \frac{|B_p \cap B_g|}{|B_p \cup B_g|}$$

Where $|B_p \cap B_g|$ is the area of the intersection and $|B_p \cup B_g|$ is the area of the union.

C Experiments

This section contains prompts and additional implementation details.



Reasoning Failure Category: Illogical Conclusion

Question: Is the difference between blue and green graph in 2008 the maximum among years?

Reasoning generated from GPT-4o: In 2008, the difference between the blue line (No) and the green line (Yes) is 43% (69% - 26%). Examining other years, the differences are: 58% in 2009, 48% in 2017, and 57% in 2018. The difference in 2008 is not the maximum among these years, with the maximum difference occurring in 2009.

Explanation: The years and numbers were incorrectly identified. An illogical conclusion was drawn.

Reasoning Failure Category: Color Mismatch

Question: Is the sum of highest and lowest value of orange graph greater than 60?

Reasoning generated from GPT-4o: The highest value of the orange graph (Russia favorability in U.S.) is 49 in 2010, and the lowest value is 15 in 2015. Adding these two values gives 49 + 15, which equals 64. Since 64 is greater than 60, the answer is 'Yes'.

Explanation: Reasoning generated from GPT-4o for the orange graph contains an inaccuracy - it identified the lowest point as 15 in 2015 which is lowest for green graph; therefore its a color mismatch.

Figure 13: Examples of reasoning failures from GPT-4o illustrating two primary error categories. Left: Illogical conclusion where the model incorrectly calculates the 2008 difference as 43% (69%-26%) but then examines wrong years (2009, 2017, 2018 instead of other years) and incorrectly concludes 2009 has the maximum difference when the question specifically asks about 2008. Right: Color mismatch where the model confuses visual elements, incorrectly identifying the lowest value of the orange line (Russia favorability in U.S.) as 15 in 2015, when this value actually corresponds to the green line (U.S. favorability in Russia), demonstrating challenges in accurate chart element identification.

C.1 Computing Infrastructure Details

Our implementation uses PyTorch 2.0 and all experiments were conducted on 4 NVIDIA A100 GPUs with 80GB of memory each. The experiments were run on Amazon Elastic Compute Cloud (Amazon EC2) instances equipped with A100 Tensor Core GPUs and 400 Gbps networking capabilities. The complete experimental pipeline took approximately 100 hours.

C.2 Prompting Strategies for Attribution

We experimented with zero-shot and few-shot prompting strategies for both VQA-based and VQR-based attribution.

Zero-Shot Prompt

System Prompt: You are a helpful assistant that responds in markdown. Help me with my math question.

Input Format:

- Chart: [chart_image], Question: [question_text], Answer: [answer_text]

User Prompt: Given this chart and the question-answer pair: question = "question", answer = "answer"; ONLY generate bounding box coordinates in X1, Y1, X2, Y2 format - A list of tuples, each containing (x1, y1, x2, y2) representing the bounding box coordinates without additional text which represents which part of the chart corresponds to the answer.

Few-Shot Prompt

System Prompt: You are a helpful assistant that responds in markdown. Help me with my math question.

Example 1:

Chart: [bar_chart_image] Question: "What was the highest value in 2020?" Answer: "85 units"
Bounding Box: (120, 45, 140, 230)

User Prompt: Given this chart and the question-answer pair: question = "question", answer = "answer" and examples; ONLY generate bounding box coordinates in X1, Y1, X2, Y2 format - A list of tuples, each containing (x1, y1, x2, y2) representing the bounding box coordinates without additional text which represents which part of the chart corresponds to the answer.

VQA based Attribution For VQA-based Attribution, we used both zero shot and few shot prompting and the prompt is described in figure 14 and 15.

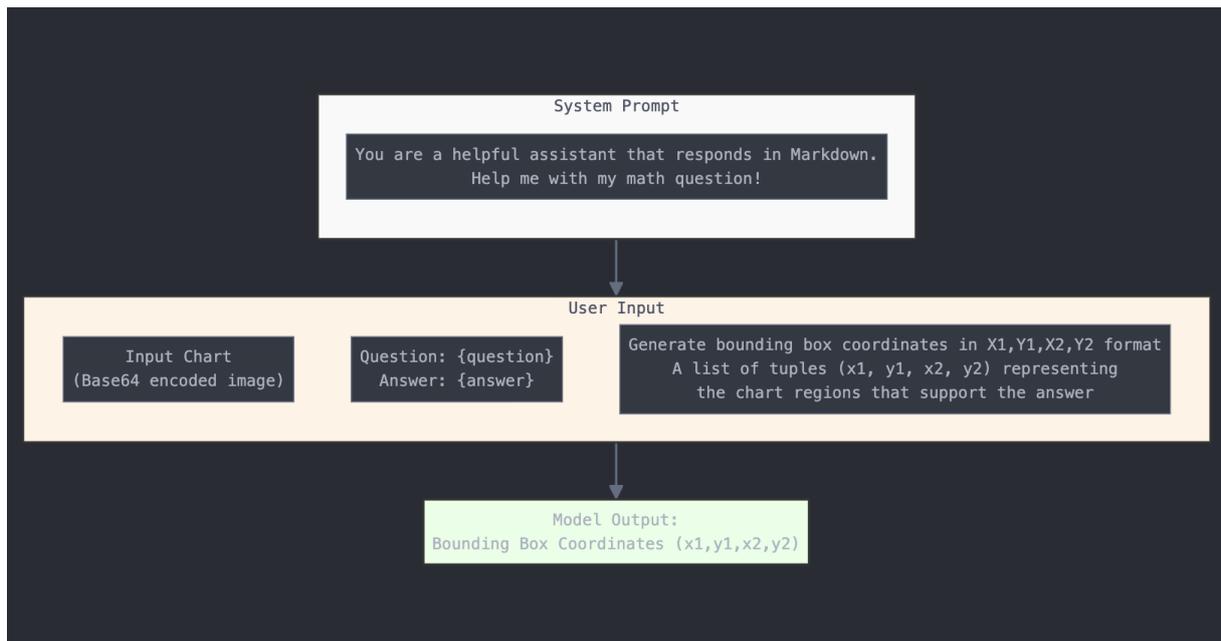


Figure 14: Zero Shot Prompting for Attribution based on VQA task.

Attribution based on VQR

For VQR based Attribution, we used both zero shot and few shot prompting and the prompt is described in figure 16 and 17.

C.3 Implementation Details - RADAR

The proposed pipeline architecture for chart understanding consists of four integral stages that work in concert to process and analyze chart images with corresponding textual inputs.

In the first stage, Input Processing, the system handles three primary inputs: the chart image, which serves as the visual input for analysis; the question and answer prompt, which guides the analysis direction. These inputs undergo Base64 encoding for the image and are formatted into a specialized text prompt structure, resulting in encoded inputs suitable for model processing.

The second stage, MLLM Processing, leverages the InternLM-XComposer2 model's multimodal capabilities to process the encoded inputs. This stage extracts Layer 16 Hidden States, which contain rich semantic information from both modalities. The image features are processed as a 35×35 patch grid, while the text features are encoded into 4096-dimensional vectors, enabling comprehensive semantic representation of both visual and textual content. This dual-stream processing ensures that both modalities contribute effectively to the final analysis.

The third stage implements a Sliding Window Attribution mechanism, which is crucial for identifying relevant regions within the chart. This process begins with window generation, where variable-sized

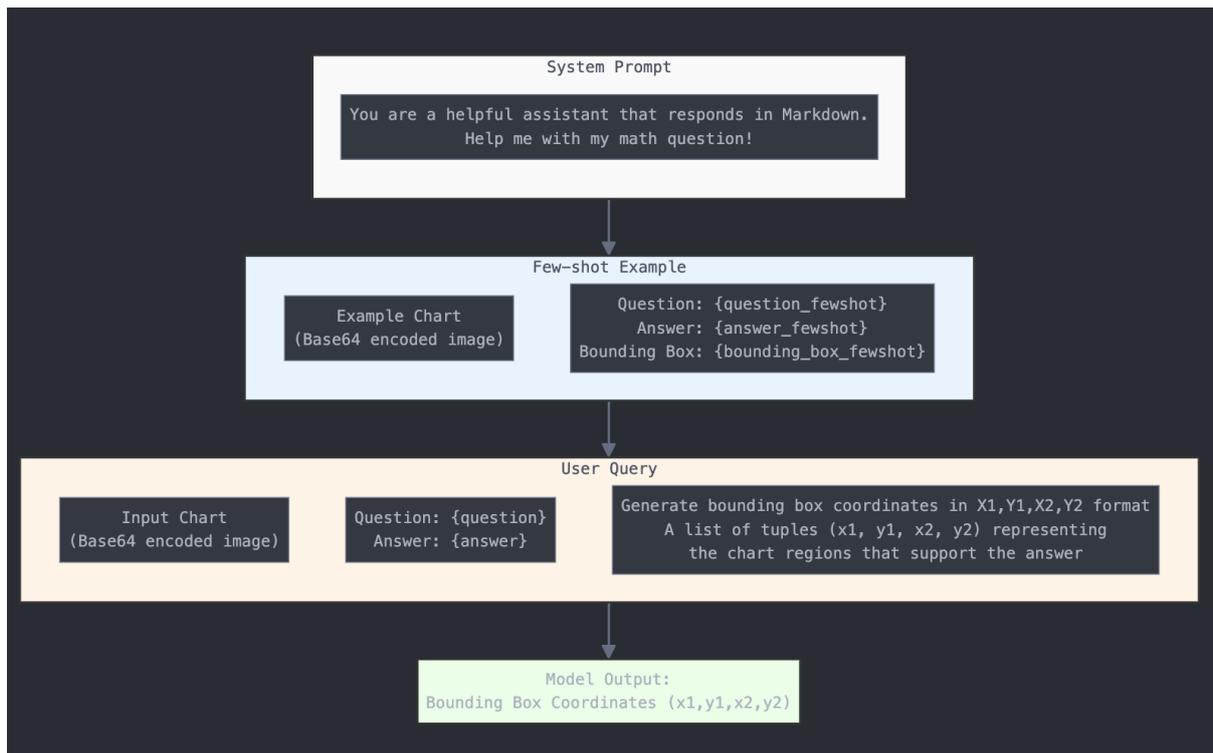


Figure 15: Few Shot Prompting for Attribution based on VQA task.

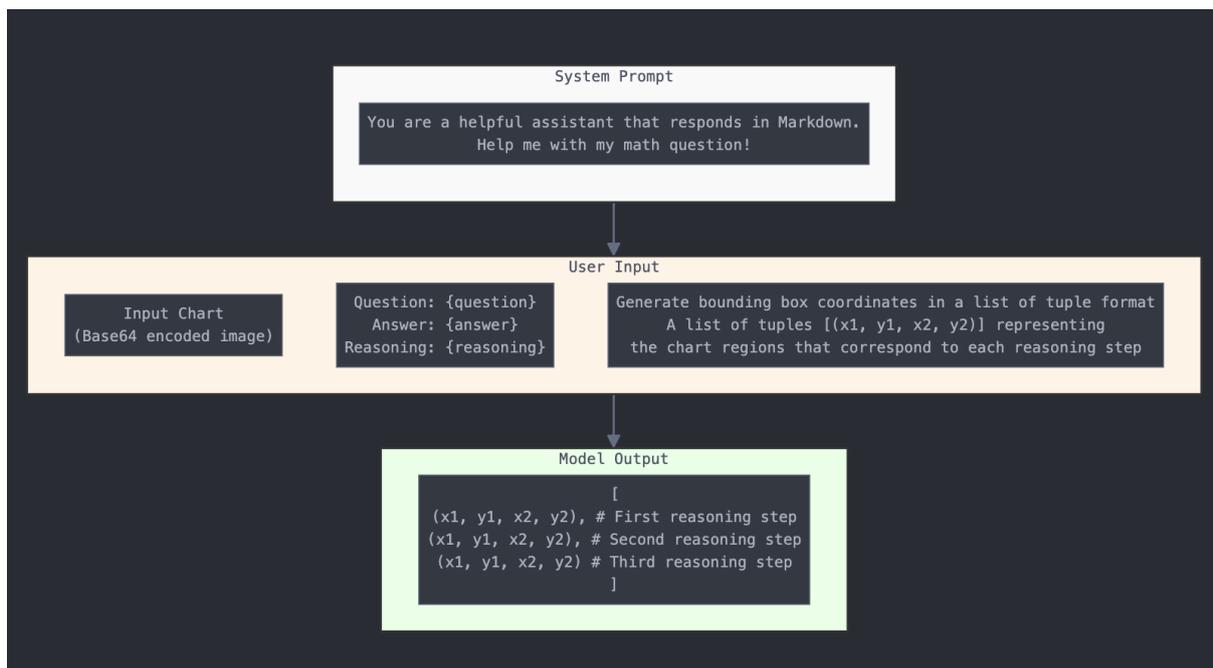


Figure 16: Zero Shot Prompting for Attribution based on VQR task.

windows are created over the image feature space. The system then computes cosine similarity between the text and image features for each window, enabling the identification of regions most pertinent to the textual input. This stage culminates in the selection of the best region, outputting coordinates (i, j, h, w) that specify both the location and dimensions of the most relevant area within the chart.

The final stage focuses on Visualization, transforming the mathematical outputs into interpretable visual representations. This involves coordinate mapping, where the model's internal coordinate space is

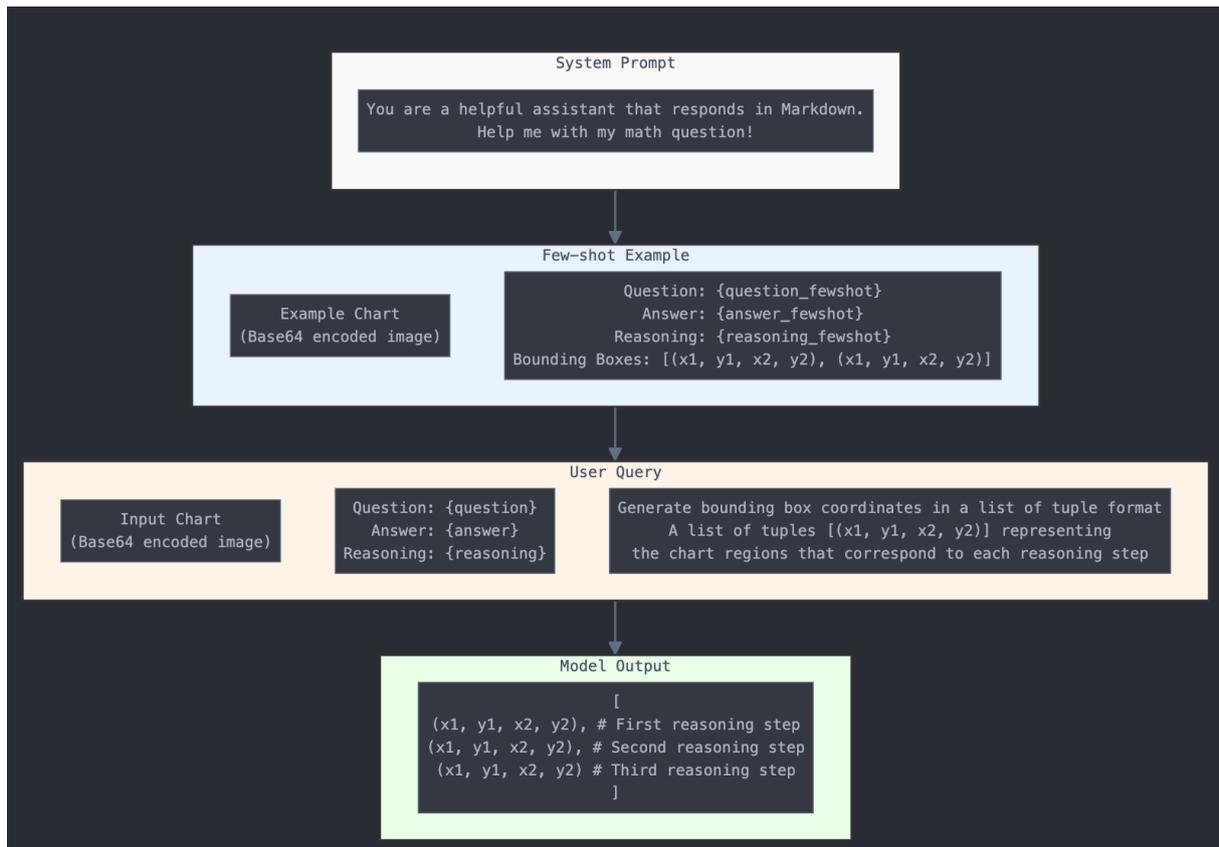


Figure 17: Few Shot Prompting for Attribution based on VQR task.

transformed into image pixel space, followed by bounding box generation that creates visible overlays highlighting the relevant regions identified by the model. This visualization stage is crucial for making the model’s decisions interpretable and useful for end users.

The entire pipeline demonstrates flexibility in handling both reasoning-based and answer-based attribution scenarios through the same architectural framework. This unified approach allows for consistent processing while accommodating different types of chart analysis tasks, from simple identification to complex reasoning about chart elements. The system maintains a consistent flow of information through each stage, ensuring that the final output effectively bridges the gap between the visual elements of the chart and the textual understanding required for comprehensive chart analysis.

Figure 18 and 19 represents VQA based and VQR based attribution details respectively.

C.4 Automatic Reasoning Evaluation

Bertscore (Zhang et al., 2019) computes the similarity between two sentences using contextual embeddings from BERT.

Semantic Textual Similarity (STS) (Agirre et al., 2012) measures how similar two pieces of text are in terms of their meaning, regardless of their exact wording.

$$STS(s_1, s_2) = \cos(\vec{v}_1, \vec{v}_2) = \frac{\vec{v}_1 \cdot \vec{v}_2}{\|\vec{v}_1\| \|\vec{v}_2\|}$$

Where \vec{v}_1 is the vector representation (embedding) of the reasoning generated from RADAR and \vec{v}_2 is the vector representation (embedding) of the reasoning that is human annotated. s_1 , and s_2 represent the reasoning generated from RADAR and human annotated ones. This equation measures the semantic similarity between RADAR-generated reasoning and human-annotated reasoning by converting both reasonings into vector representations and computing their cosine similarity using dot product and vector magnitudes.

After all the scores were calculated, we took an average and the scores are reported in table 4.

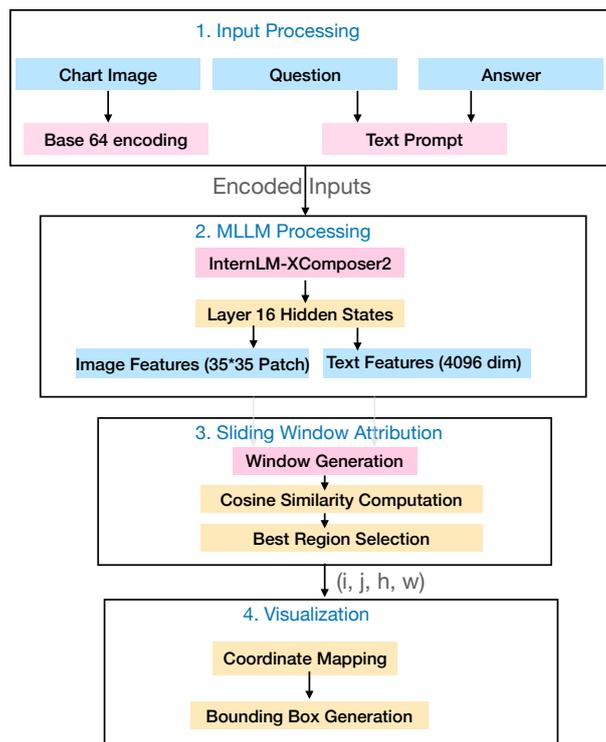


Figure 18: The pipeline architecture for chart understanding with InternLM-XComposer2 illustrates a four-stage process that bridges visual and textual modalities in chart analysis. The system progresses through Input Processing (encoding of chart images and text), MLLM Processing (multimodal feature extraction), Sliding Window Attribution (region identification), and Visualization (interpretable output generation), enabling comprehensive chart understanding through a unified architectural framework. This architecture supports answer-based attribution.

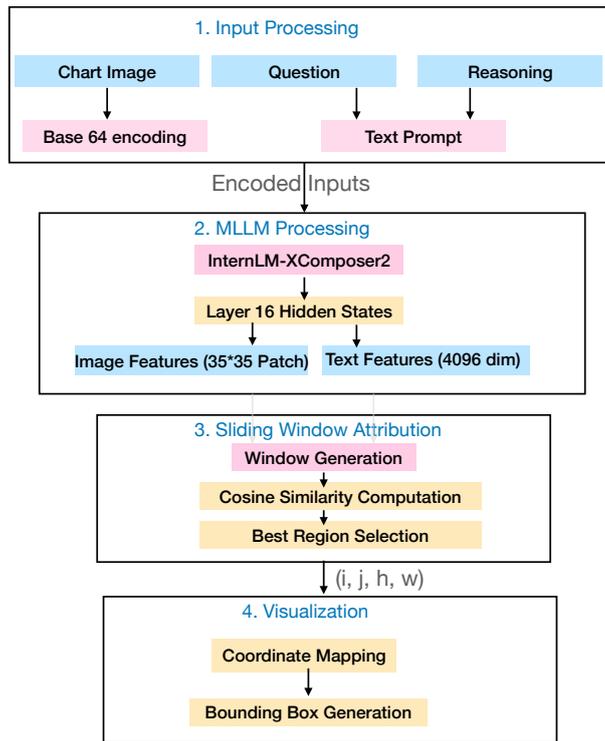


Figure 19: The pipeline architecture for chart understanding with InternLM-XComposer2 illustrates a four-stage process that bridges visual and textual modalities in chart analysis. The system progresses through Input Processing (encoding of chart images and text), MLLM Processing (multimodal feature extraction), Sliding Window Attribution (region identification), and Visualization (interpretable output generation), enabling comprehensive chart understanding through a unified architectural framework. This architecture supports reasoning based attribution.

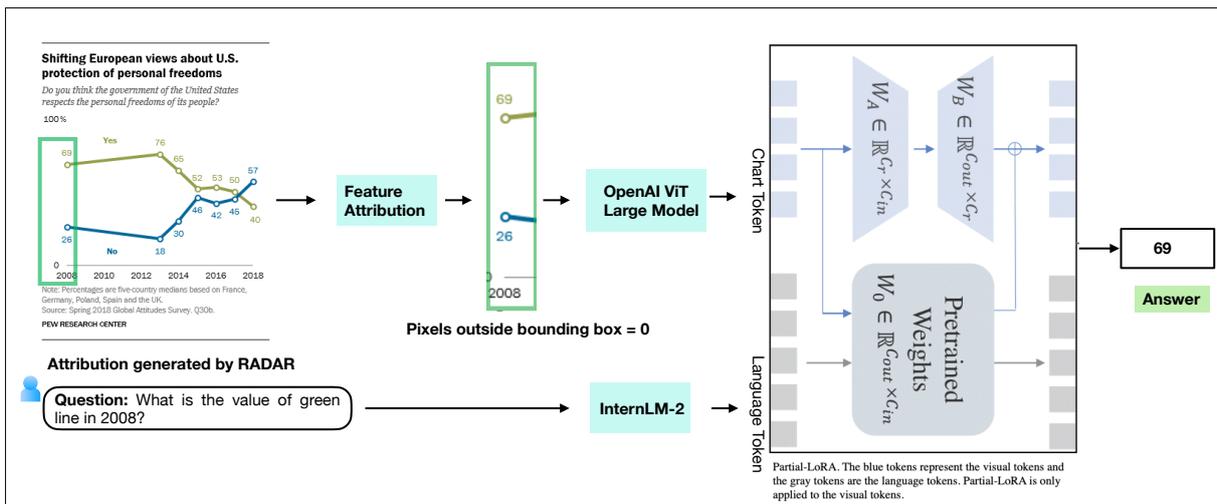


Figure 20: Automatic Attribution Evaluation: Attribution generated by RADAR is passed through feature attribution that sets the pixels outside the bounding box to zero. Feature Attributed Charts and Question is passed as input through the framework to collect answers. The generated answer is then matched with the ground truth.