

FCMR: Robust Evaluation of Financial Cross-Modal Multi-Hop Reasoning

Seunghee Kim[†], Changhyeon Kim[†], Taeuk Kim^{*}

Hanyang University, Seoul, Republic of Korea

{gyg9325, livex, kimtaeuk}@hanyang.ac.kr

Abstract

Real-world decision-making often requires integrating and reasoning over information from multiple modalities. While recent multimodal large language models (MLLMs) have shown promise in such tasks, their ability to perform multi-hop reasoning across diverse sources remains insufficiently evaluated. Existing benchmarks, such as MMQA, face challenges due to (1) data contamination and (2) a lack of complex queries that necessitate operations across more than two modalities, hindering accurate performance assessment. To address this, we present Financial Cross-Modal Multi-Hop Reasoning (FCMR), a benchmark created to analyze the reasoning capabilities of MLLMs by urging them to combine information from textual reports, tables, and charts within the financial domain. FCMR is categorized into three difficulty levels—Easy, Medium, and Hard—facilitating a step-by-step evaluation. In particular, problems at the Hard level require precise cross-modal three-hop reasoning and are designed to prevent the disregard of any modality. Experiments on this new benchmark reveal that even state-of-the-art MLLMs struggle, with the best-performing model (Claude 3.5 Sonnet) achieving only 30.4% accuracy on the most challenging tier. We also conduct analysis to provide insights into the inner workings of the models, including the discovery of a critical bottleneck in the information retrieval phase.

1 Introduction

Despite the recent progress in AI (Touvron et al., 2023; Anthropic, 2024; OpenAI, 2024a), developing systems capable of human-level reasoning remains a challenge. Human cognition involves integrating information from multiple modalities to comprehend and make decisions. One domain that requires such a comprehensive understanding

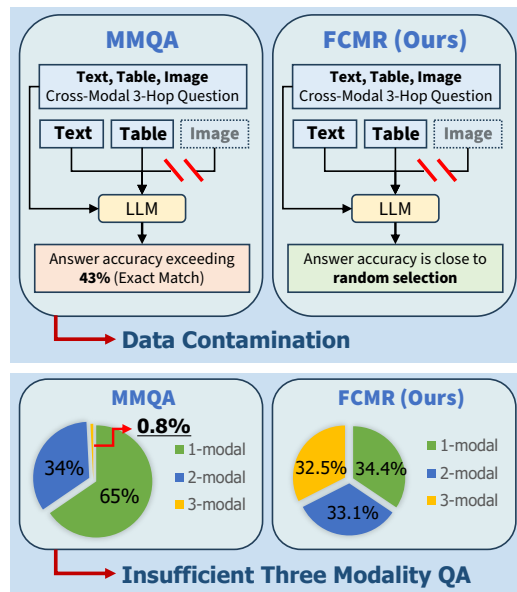


Figure 1: Prior benchmarks for cross-modal multi-hop reasoning, such as MMQA (Talmor et al., 2021), exhibit some flaws. MMQA’s cross-modal three-hop questions are often solvable without images, and its complexity is limited, with only 0.8% of instances having three modalities. In contrast, FCMR addresses these issues.

is finance (Xie et al., 2024), where analysts must simultaneously examine textual reports, tabular data (e.g., balance sheets), and visual data (e.g., charts). For example, verifying the statement from Figure 2—“The corporation, with the smallest act value in the years when the fopo value for ABBOTT LABORATORIES is below 730.5, is entitled to receive \$43 million in minimum sublease income from non-cancelable subleases.”—one must consider all the clues provided by each source, an ability we refer to as **cross-modal multi-hop reasoning**.

While the literature (Chen et al., 2020; Hanan et al., 2020; Talmor et al., 2021; Chang et al., 2022) presents initial attempts to evaluate the cross-modal multi-hop reasoning capabilities of multimodal large language models (MLLMs), these efforts exhibit several critical shortcomings that un-

[†]Equal contribution. ^{*}Corresponding author.

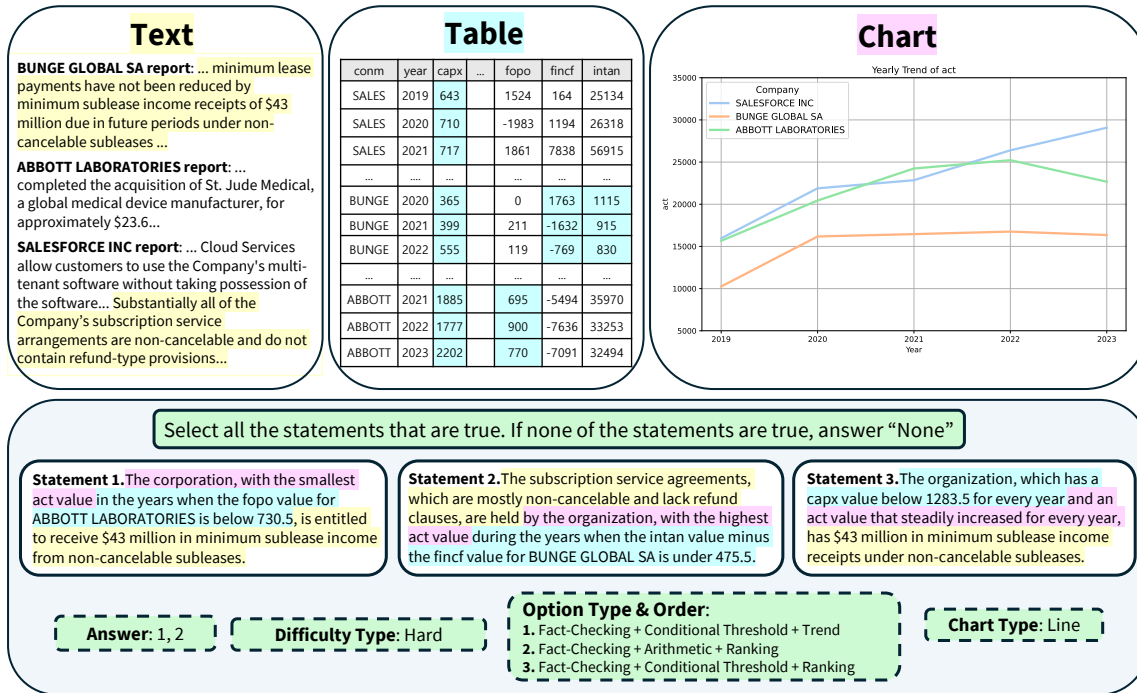


Figure 2: An example from FCMR at the Hard difficulty level, where all statements require cross-modal three-hop reasoning. Highlights in yellow, cyan, and pink denote information from text, tables, and charts, respectively. The model must list all true statements and is correct only if its final prediction (“1, 2” in this case) is accurate. Information within dashed lines is used only for data generation and excluded from actual instances.

dermine their robustness. First, the heavy reliance on Wikipedia as the foundation for most benchmarks raises concerns. As Wikipedia is widely known to be a key resource in the pretraining of many recent models, evaluations using Wikipedia-based datasets risk introducing inherent biases. These biases may skew results in favor of models that simply recall memorized knowledge, rather than accurately assessing reasoning abilities on unseen data. Moreover, the scope of validation needs to expand to encompass professional domains, such as finance and science.

Second, current benchmarks are largely focused on testing straightforward problems, such as single- and two-hop reasoning. As shown in Figure 1, MMQA (Talmor et al., 2021)—one of the leading benchmarks in this field—features only about 0.8% of the queries that explicitly require *three-hop* cross-modal reasoning. Furthermore, in preliminary experiments, we discovered that GPT-4o (OpenAI, 2024a) can solve the MMQA’s most challenging problems with a 43% exact match accuracy, even without access to visual clues. This result highlights the urgent need to establish a higher standard for evaluating cross-modal multi-hop understanding in a more robust and reliable manner.

In this work, we propose **Financial Cross-**

Modal Multi-Hop Reasoning (FCMR),¹ a novel benchmark designed to address the limitations of existing datasets in cross-modal multi-hop reasoning. FCMR provides multiple-choice QA samples that test the integration of facts from text, tables, and charts. It consists of three levels of difficulty: Easy, Medium, and Hard. As in Figure 1, all instances in FCMR necessitate understanding three modalities to be answered correctly. In addition, problems at the Hard level explicitly demand cross-modal three-hop reasoning, making them more challenging (Figure 2). As FCMR is built using data sources from the financial domain, it is relatively free from the risk of data contamination.

Experiments on FCMR confirm that it poses challenges even for state-of-the-art MLLMs, e.g., GPT-4o and Claude 3.5 Sonnet, encouraging research efforts to develop systems capable of reasoning across multiple modalities. For analysis, we define four procedures of cross-modal multi-hop reasoning—Planning, Modality Identification, Information Retrieval, and Information Reasoning—and probe diverse models. We reveal that models particularly struggle with the Information Retrieval phase, implying that MLLMs often fail to extract

¹Our dataset is available at <https://github.com/HYU-NLP/FCMR>.

Benchmarks	Cross-Modal 2-Hop?	Cross-Modal 3-Hop?	Contain Table?	Contain Image?	Domain Specific?
ManyModalQA	✗	✗	✓	✓	✗
CT2C-QA	✗	✗	✓	✓	✓
MME-Finance	✗	✗	✓	✓	✓
WebQA	✓	✗	✗	✓	✗
MuMuQA	✓	✗	✗	✓	✓
FinQA	✓	✗	✓	✗	✓
TAT-QA	✓	✗	✓	✗	✓
HybridQA	✓	✗	✓	✗	✗
OTT-QA	✓	✗	✓	✗	✗
TANQ	✓	✗	✓	✗	✗
MMQA	✓	✗	✓	✓	✗
FCMR (Ours)	✓	✓	✓	✓	✓

Table 1: Comparison of datasets based on cross-modal reasoning and modality coverage. MMQA’s ✗ shows that although it includes cross-modal three-hop reasoning, such instances constitute only 0.8% of the dataset.

information from a specific modality, even when they successfully identify where the required facts are located. We also report findings from additional analyses, including the observation that MLLMs have difficulty counting negative numbers.

2 Related Work

2.1 Cross-Modal Multi-Hop Reasoning

Benchmarking cross-modal multi-hop reasoning has received considerable attention. Efforts include WebQA (Chang et al., 2022) and MuMuQA (Reddy et al., 2022)—for two-hop cross-modal reasoning with text and images—as well as HybridQA (Chen et al., 2020), OTT-QA (Chen et al., 2021a), FinQA (Chen et al., 2021b), TAT-QA (Zhu et al., 2021), and TANQ (Akhtar et al., 2024), which include text and tables. However, all these datasets are limited to **only two modalities**, making them inadequate for evaluating more complex cases.

Meanwhile, datasets like ManyModalQA (Hanan et al., 2020), CT2C-QA (Zhao et al., 2024), and MME-Finance (Gan et al., 2024) incorporate three modalities but lack an inherent focus on cross-modal multi-hop reasoning.

MMQA (Talmor et al., 2021), in contrast, deals with **three modalities**—text, tables, and images—and requires **three-hop reasoning**, setting it apart from others. It has served as the de facto standard for evaluating related methods (Rajabzadeh et al., 2023; Yu et al., 2023a; Luo et al., 2023; Zhang et al., 2024a; Abaskohi et al., 2024). MMCQA (Li et al., 2022) and MMCV (Wang et al., 2024) have been developed as extensions of MMQA.

Dataset: MMQA	Image?	Exact Match (%)	F1 Score (%)
Random Selection	-	0.0	1.2
GPT-4o	✗ ✓	43.4 63.4	46.2 67.5

Table 2: Experiments on a subset of MMQA requiring cross-modal three-hop reasoning reveal that GPT-4o performs reasonably well even without images. This implies that it already contains information derivable from input images, questioning the rigor of MMQA. For more details, refer to Appendix C.2.

The characteristics of the benchmarks are summarized in Table 1. FCMR is crafted to address the limitations of previous ones, particularly MMQA.

2.2 Limitations of MMQA

We briefly revisit the drawbacks MMQA to emphasize the need for a new, robust benchmark for cross-modal multi-hop reasoning.

Data Contamination As MMQA is constructed from Wikipedia, it is vulnerable to data contamination. That is, the model being tested may already possess internalized knowledge of certain facts, reducing its dependence on the dataset’s provided input. In particular, Table 2 shows that GPT-4o can achieve reasonable performance on the most challenging part of MMQA—questions intentionally tailored for requiring a combination of information from three modalities—*without* relying on visual hints.² This suggests that MMQA falls short of effectively measuring cross-modal multi-hop reasoning ability as it was originally intended.

Lack of Cross-Modal Three-Hop Cases Only about 0.8% (205 instances) of the MMQA dataset consists of cross-modal three-hop reasoning, while the majority comprises either one-hop or two-hop questions. This scarcity restricts its effectiveness in thoroughly evaluating a model’s performance on complex reasoning tasks with intricate interactions across text, tables, and images.

3 Proposed Benchmark: FCMR

We introduce **Financial Cross-Modal Multi-Hop Reasoning (FCMR)**, a new benchmark created to alleviate the shortcomings of MMQA and enable a more comprehensive evaluation of cross-modal multi-hop reasoning. FCMR includes three modalities—text, tables, and charts—and presents questions that entail selecting all correct statements

²Figure 10 presents an example of data leak in MMQA.

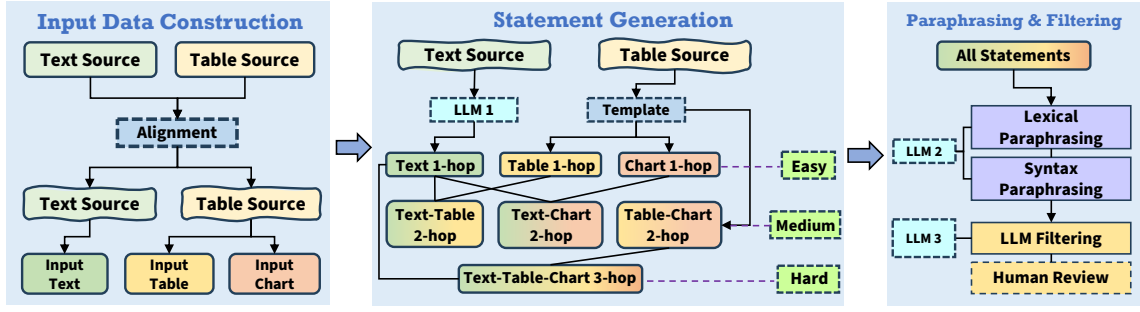


Figure 3: **CMRGen** is an automated and efficient framework for building cross-modal multi-hop reasoning datasets. (1) **Modality Data Construction** extracts text, table, and chart modalities from sources sharing common entities. (2) **Statement Generation** produces cross-modal one-, two-, and three-hop statements using LLM and templates. (3) **Paraphrasing & Filtering** applies two-stage paraphrasing with LLMs, followed by LLM & Human filtering.

from a set of three. The tested model must identify all true statements in the problem and is considered correct only if its final prediction is accurate. This task inherently requires the model to address numerous subtasks—e.g., numerical calculation and chart interpretation—that require deep reasoning.

3.1 Dataset Generation Framework: CMRGen

We propose **Cross-Modal Multi-Hop Reasoning Generator (CMRGen)**, a framework that facilitates the construction of cross-modal multi-hop reasoning datasets across various domains. CMRGen distinguishes itself from other cross-modal multi-hop dataset generation methods with its highly automated and cost-effective pipeline. Notably, while generating a single question in MMQA incurs a cost of \$0.33, our method reduces this to \$0.004 per question. Furthermore, the framework demonstrates high adaptability to various domains and offers seamless control over difficulty levels, ranging from Easy to Hard. In this study, we focus on the financial domain, where complex interactions among text, tables, and charts frequently occur, making it an ideal testbed for evaluating cross-modal multi-hop reasoning. However, the proposed pipeline is also readily applicable to domains other than finance, demonstrating its flexibility.³

3.2 Procedure of CMRGen

CMRGen have three stages, as depicted in Figure 3: (1) Input Data Construction, (2) Statement Generation, and (3) Paraphrasing & Filtering. We explain each step using **FCMR** as an example. Details of the procedure can be found in Appendix B.

(1) Input Data Construction In the first stage, we prepare and organize data for the text, table, and chart modalities. As the origin of information, CMRGen utilizes two sources: Text Source and Table Source. For FCMR, the Text Source consists of Annual 10-K Reports collected from the SEC EDGAR database,⁴ while the Table Source is derived from Annual Simplified Financial Statements provided by WRDS Compustat.⁵ We then filter entries that share common company entities, aligning the two sources. Finally, we construct each data instance in FCMR, comprising a document, a table, and a chart about three companies. In the next step, this instance will be supplemented with three statements serving as questions about the contents generated in this stage. Note that the chart is created by plotting specific columns from the Table Source using custom scripts, and the columns used to create the chart are removed from the table. We create synthetic charts but ensure diversity by incorporating various chart types and styles, as well as using different libraries (see Appendix B.6).

(2) Statement Generation In the second phase, diverse forms of statements (i.e., questions) are crafted for each FCMR instance. We leverage GPT-4o-mini to generate text-based one-hop statements by extracting relevant facts from the Text Source. In addition, by leveraging various templates tailored to real-world financial scenarios—such as Trend, Ranking, Conditional Threshold, and Arithmetic—we create table-based and chart-based one-hop statements based on the Table Source. We then combine these single-modal one-hop statements across entities to construct cross-modal two-

³We showcase the application of the proposed method in material science. For more details, refer to Appendix A.

⁴<https://www.sec.gov/search-filings/>

⁵<https://wrds-www.wharton.upenn.edu/pages/grid-items/compustat-global-wrds-basics/>

hop statements, which are further merged to create cross-modal three-hop statements. Each statement is categorized into Easy, Medium, or Hard based on the number of hops required for reasoning. The complete taxonomies of statement types and templates are presented in Table 10 and Table 11.

(3) Paraphrasing & Filtering Lastly, we apply two stages of lexical and syntactic paraphrasing using GPT-4o to enhance the diversity of expressions in statements. Afterward, we conduct LLM-based filtering, using Claude 3.5 Sonnet, to ensure semantic accuracy. By employing separate LLMs in each step, we aim to mitigate model-oriented biases. Our framework also allows for an optional human expert review to improve their quality further. For FCMR, we use this review step to ensure high standards, especially for challenging (i.e., Hard-level) problems. Appendix C.1 provides the specific details of this human annotation process.

3.3 Distractor Generation

Instead of simply adjusting numerical values to generate incorrect statements, we reflect realistic scenarios in the financial domain, where analysis of multiple companies is common, by generating distractors based on corporate entities. Since each statement in Easy, Medium, and Hard levels is combined with corporate entities, we generate distractors by replacing them with other companies in the same instance.

3.4 Multiple-Choice Design

Previous research on cross-modal multi-hop reasoning has often employed descriptive or short-answer formats, evaluated with metrics like F1 or Exact Match. These approaches might result in inaccurate evaluations by penalizing semantically appropriate answers that slightly deviate in form. To address this, we adopt a multiple-choice format with three statements. Recent work (Pang et al., 2024) argues that single-choice question formats are more suitable for model evaluation than free-form answers, supporting our decision. In contrast to existing multimodal benchmarks (Li et al., 2023; Yue et al., 2024; Ying et al., 2024; Zhang et al., 2024b; Liu et al., 2024b) that typically rely on a single correct answer, our setup allows for zero to three correct statements. This strategy enhances the complexity of the reasoning process needed to answer the problem accurately, requiring a more comprehensive synthesis of all provided

Dataset	WPD	LD
MRPC	0.12	0.42
PAWS	0.07	0.13
FCMR (Ours)	0.2	0.45

Table 3: WPD (Word Position Deviation) represents syntactic diversity, and LD (Lexical Deviation) reflects lexical diversity. Both metrics indicate that higher scores correspond to greater diversity in paraphrasing.

Dataset: FCMR (Hard)	Image?	Accuracy
Random Selection	-	12.28
GPT-4o	✗ ✓	14.71 24.37

Table 4: Replication of the experiments from Table 2 with FCMR. GPT-4o’s performance drops to near random selection when charts are omitted, suggesting that FCMR is relatively robust against data contamination.

statements. It also allows for precise evaluation of models’ cross-modal multi-hop reasoning abilities.

3.5 Data Quality Control

To uphold high data quality, we implement multifaceted verification protocols.⁶ Specifically, we utilize Word Position Deviation (WPD) and Lexical Deviation (LD) metrics (Liu and Soh, 2022) to evaluate paraphrasing quality and compare these values with those from MRPC (Dolan and Brockett, 2005) and PAWS (Zhang et al., 2019). The outcomes, presented in Table 3, confirm the superiority of our paraphrasing method. Furthermore, to verify that our dataset avoids the contamination issue identified in MMQA, we replicate the contamination experiment under the same conditions. As shown in Table 4, when the chart images are withheld, GPT-4o’s performance approximates random guessing, alleviating the risk of data contamination in FCMR. We further mitigate potential biases by equalizing modality order, statement types, and answer distributions. Figure 21 demonstrates that our benchmark is well-balanced across various perspectives. The final dataset includes 757 Easy, 728 Medium, and 714 Hard instances. Hard-level problems demand precise cross-modal three-hop reasoning. While statements in Easy and Medium questions involve simpler reasoning (e.g., single-modal one-hop for Easy), resolving each instance still necessitates processing information from all three provided modalities—text, table, and chart.

⁶Refer to Appendix C for full details of our strategies.

Metric: Accuracy (%)	Easy	Medium	Hard	Avg
Random Selection	12.2	12.91	12.28	12.46
Multimodal Large Language Models (MLLMs)				
ChartInstruct-Llama2	11.49	12.64	10.78	11.64
llama3-llava-next-8b-hf	16.86	12.22	11.53	13.54
MiniCPM-V-2_6	16.38	11.68	13.17	13.74
Qwen2-VL-7B-Instruct	17.57	13.32	12.04	12.32
Llama 3.2 90B-Vision	42.47	21.60	13.73	25.94
GPT-4o mini	49.14	21.98	13.03	28.05
Gemini 1.5 Flash	57.33	26.65	13.43	32.80
Gemini 1.5 Pro	63.01	31.18	22.27	38.82
GPT-4o	64.20	43.70	24.37	44.09
Claude 3.5 Sonnet	75.43	50.82	30.39	52.21
Large Language Models (LLMs) with Deplot				
Qwen2-7B-Instruct	21.66	11.95	14.01	15.87
Llama 3.1 8B-Instruct	30.91	13.05	10.36	18.11
Llama 3.1 70B-Instruct	46.37	17.86	14.01	26.08
Llama 3.2 90B-Vision	50.20	22.39	11.90	28.16
GPT-4o mini	57.60	26.51	12.61	32.24
GPT-4o	68.69	49.18	32.91	50.26
Claude 3.5 Sonnet	66.84	46.15	36.13	49.71

Table 5: Performance of MLLMs and LLMs on FCMR. For LLMs, charts are converted into tables using Deplot. The best performance at each difficulty level and category is highlighted in **bold**.

3.6 Chart Type Coverage

To ensure sufficient diversity in chart types, FCMR includes four commonly used formats in the financial domain: line, bar, scatter, and pie charts. These chart types reflect the typical visualizations found in real-world financial documents. A recent analysis (Christensen et al., 2024) of 23,906 quantitative infographics extracted from all 10-K reports filed with the SEC between 2003 and 2019 shows that bar charts account for 55.0%, pie charts for 30.6%, and line charts for 12.7%. The chart types included in our FCMR dataset directly correspond to these prevalent quantitative categories, collectively covering approximately 98% of the most frequently used quantitative visualization types in actual 10-K filings. Details on the chart generation process are provided in Appendix B.6.

4 Experiments

4.1 Experimental Setup

We evaluate a wide range of MLLMs and LLMs on FCMR under a zero-shot CoT setting, where no task-specific tuning or demonstration is provided. This setting reflects common use cases for MLLMs, ensuring an unbiased model evaluation and capturing overall performance trends. All models are prompted with the same template in Figure 23. Tables are represented in JSON format. For

proprietary models, we employ the Claude version *claude-3-5-sonnet-20241022*, the GPT-4o version *gpt-4o-2024-08-06*, and Gemini version *gemini-1.5-pro-002*. We also test several open-source models: Llama variants (Touvron et al., 2023), Qwen (Yang and Yang, 2024), MiniCPM (Yao and Yu, 2024), Llava (Liu et al., 2024a), and ChartInstruct (Masry and Shahmohammadi, 2024).

4.2 Main Results

Performance of MLLMs Table 5 reports the performance of various MLLM across different levels. Most open-source models perform just above random chance at the Easy level, focused on single-modal, one-hop reasoning, confirming FCMR as a challenging benchmark. Proprietary models perform better, demonstrating a remarkable gap in reasoning ability. However, at the Hard level, which necessitates full cross-modal three-hop integration, even sophisticated models, including Claude 3.5 Sonnet, achieve only around 30%.⁷ This result underscores FCMR’s challenging nature and the need for developing more advanced reasoning strategies.

Performance of (M)LLMs + Deplot For image-blind standard large language models (LLMs), we use Deplot (Liu et al., 2023) to convert charts into tables, ensuring that all models can access chart information. We also explore applying the same heuristic to MLLMs, as the literature suggests that MLLMs tend to rely more on textual clues than visual ones (Rahmanzadehgervi et al., 2024).

Experimental results indicate that open-source models with fewer than 8B parameters continue to perform comparably to random selection for tasks at the Medium and Hard levels. However, for the Easy category, they demonstrate superior performance compared to similarly sized MLLMs. Surprisingly, even advanced MLLMs such as GPT-4o and Claude 3.5 Sonnet achieve performance gains in certain cases, suggesting that their visual interpretation capabilities are still not perfect. In Section 6.2, we dive deeper into this phenomenon.

5 Analysis

In this section, we analyze the inner workings of closed-source MLLMs—GPT-4o, Claude 3.5 Sonnet, and Gemini 1.5 Pro—which outperform

⁷Recently, deep reasoning models, e.g., OpenAI’s o1 (OpenAI, 2024b), have emerged. While they are not directly related to multimodal functionality, we also evaluate them at the Hard level. As a result, we confirm that o1 achieves 43% accuracy, indicating they are still far from perfect (refer to Table 9).




Figure 4: Fine-grained stage-based analysis of three advanced MLLMs. This unique strategy discovers several intriguing findings, including that the models fail most often at (3) Information Retrieval.

		Ground Truth	
		True	False
Prediction	True	514	178
	False	557	893
Claude 3.5 Sonnet		Precision 74.28	Recall 47.99
		F1-score 58.31	Accuracy 65.69

		Ground Truth	
		True	False
Prediction	True	455	265
	False	616	806
GPT-4o		Precision 63.19	Recall 42.48
		F1-score 50.81	Accuracy 58.87

		Ground Truth	
		True	False
Prediction	True	470	326
	False	604	745
Gemini 1.5 pro		Precision 59.05	Recall 43.88
		F1-score 50.35	Accuracy 56.72

Figure 5: Confusion matrices for three advanced MLLMs, with metrics in percentages (%).

smaller open-source models. We focus on their performance at the Hard level, as this subset presents the most challenging questions for the models.

5.1 Statement-Level Analysis

While each problem in FCMR requires models to draw an overall conclusion on three statements, their partial solutions for each statement can provide insight into how well each model handles diverse cases. We gather statistics on each model’s predictions for every statement and construction confusion matrices for analysis. We have $714 \times 3 = 2,142$ statements with gold-standard answers, evenly split into true (1,071) and false (1,071) ones. Each model’s prediction is annotated for these statements, forming the matrices in Figure 5.

While all three models show limitations in precision, recall, and F1-score, Claude 3.5 Sonnet achieves comparatively better performance. With a high precision of 74.27, Claude effectively mini-

mizes false positives, reflecting an ability to accurately classify positive cases. However, its recall remains limited to 47.99, indicating a reduced capacity to capture all true positives. Despite this trade-off, Claude achieves the highest accuracy at 65.69 and an F1-score of 58.31, outperforming the other two models, GPT-4o and Gemini 1.5 Pro.

Moreover, all three models display a notable tendency to adopt a conservative decision-making strategy by defaulting to *false* in cases of uncertainty or low confidence. This behavior reflects a low-risk approach aimed at reducing false positive classifications, even if it results in a lower recall.

5.2 Stage-Based Analysis

In the main experiments, we observed advanced MLLMs follow a similar sequence of reasoning steps to solve problems in FCMR. Based on this, we define four fine-grained reasoning steps to identify where errors commonly occur. The four stages are specified as: **(1) Planning**: identifying the required values, **(2) Modality Identification**: recognizing which modality contains these values, **(3) Information Retrieval**: extracting relevant information from the identified modality, and **(4) Information Reasoning**: reasoning over the extracted information under the given conditions. Each instance includes three statements, each requiring the four-step process across all three modalities. Models must execute steps (1)-(4) three times per statement and repeat this for all three statements before answering (see Figure 4 (a)).

Level / Modality	Text	Table	Chart	Total
Easy	1 (4%)	5 (21%)	18 (75%)	24
Medium	5 (16%)	6 (19%)	20 (65%)	31
Hard	6 (14%)	13 (32%)	22 (54%)	41

Table 6: Error counts and proportions by modality for Claude 3.5 Sonnet across 90 statements per level.

We manually monitor MLLMs’ inference trajectories on 40 given samples. After each fine-grained stage, we record the number of problems successfully processed by the models, forming a success history diagram in Figure 4 (b). The visualization reveals an intriguing pattern: for most samples, MLLMs fail at some stage before completing the reasoning steps for the first statement. Specifically, we observe a notable performance drop at the [Statement 1, Modality 2, (3)-(4)] stage. This suggests that while the models handle the first modality relatively well, they struggle considerably when they encounter a second modality. Interestingly, GPT-4o outperforms Claude 3.5 Sonnet in the first modality reasoning step of the first statement, but Claude surpasses GPT-4o starting from the second modality phase. As the models progress to the second statement, Claude’s performance diverges further from the others, showcasing more robust and sustained reasoning capabilities.⁸

We further explore model failures by identifying which of the four reasoning steps (1)-(4) these failures occur in, regardless of statement and modality. As illustrated in Figure 22, the most common cause of failure across MLLMs occurs at step (3) Information Retrieval—failing to extract the required information from the identified modality. The second most frequent failures arise at stage (4), Information Reasoning, where models struggle to correctly apply the retrieved information to the given conditions. Notably, Gemini 1.5 Pro exhibits a higher proportion of failures at step (4), meaning that even when it successfully retrieves information, it has difficulty reasoning over it.⁹

While Claude and GPT make no modality identification errors (stage (2)), Gemini 1.5 Pro occasionally misidentifies modalities, such as confusing chart values with table values. Gemini 1.5 Pro has no failures at the Planning stage (stage (1)). In contrast, GPT and Claude sometimes skip planning for the third modality after successfully handling the first two, leading to task failure. This shows that

⁸Fine-grained model answer examples are in Figure 11.

⁹Examples are shown in Figure 12 and Figure 13.

Level / Type	Line	Bar	Scatter	Pie	Total
Easy	74.89	78.60	71.01	84.31	75.43
Medium	52.70	50.00	49.79	-	50.82
Hard	39.22	29.20	23.44	-	30.39

Table 7: Accuracy by chart type, based on Claude 3.5 Sonnet. All values are presented as percentages (%). Pie charts are only used in the Ranking option type of the Easy difficulty, as they are unsuitable for the Trend option type.

while all models struggle at later reasoning steps, GPT and Claude particularly struggle to maintain a consistent strategy across modalities.

6 Case Study

Given Claude 3.5 Sonnet’s effectiveness for FCMR, we conduct case studies to derive insights for enhancing cross-modal multi-hop reasoning.

6.1 Error Rate by Modality

Table 6 displays the numbers and proportions of statements Claude fails to interpret correctly, based on random 90 statements for each difficulty level.¹⁰ At the Easy level, Claude often struggles with analyzing charts, performing noticeably worse compared to its handling of text and tables. This disparity indicates that MLLMs exhibit weaker capabilities in interpreting charts than in processing textual or tabular data (see Appendix 6.2 for details). However, as the difficulty level increases, errors in text and tables become more prominent. This shift is likely attributed to the increased complexity of reasoning chains, which raises the likelihood of errors in processing text and tables.

6.2 Chart Interpretability

Building on previous findings that even Claude struggles with chart interpretation, we analyze the specific conditions that pose the greatest challenges. Table 7 shows that among line, bar, scatter, and pie charts, scatter plots are the most challenging due to their less structured representations. In contrast, MLLMs find it easier to identify trends in line or bar charts, which provide clearer patterns. Ranking tasks appear simpler than trend analysis, as they involve identifying extremes, whereas trend detection demands more advanced inference.

¹⁰Errors unrelated to modality, such as misinterpreting conditions, are excluded from this analysis.




Figure 6: Error categories and their distributions across 100 manually labeled error cases for Claude 3.5 Sonnet.

6.3 Error Classification and Inspection

To manually inspect Claude’s working patterns, we examine 100 error cases where it was unable to provide accurate answers. The distribution of these errors is visualized in Figure 6. Real error cases can be found in Figure 7 and Appendix D.

The most common issue involves misinterpreting trends in charts (35 cases). The second most frequent error type (16 samples) pertains to the misidentification of top-ranked entities or overall rankings. The model fails in 11 cases to correctly identify entities that meet the given conditions or misapplies the conditions. 17 errors arise from misinterpreting data, including sums, negative values, subtle differences, and column confusion. 4 involve conflating company or entity identities. The model also makes 8 errors in information extraction and reasoning, such as misreading facts or drawing unjustified conclusions. Lastly, 9 cases stem from misinterpreting instructions, ignoring required modalities, or making illogical inferences.

These results highlight that successful cross-modal multi-hop reasoning demands both strong reasoning abilities and effective interpretative skills for each modality. Therefore, enhancing a model’s capacity to interpret individual modalities is essential for enabling effective multi-hop reasoning.

6.4 Preliminary Task Optimization

To explore ways to enhance MLLM performance on FCMR, we conduct preliminary experiments in modality integration and prompt optimization. Table 8 presents the results of applying three techniques—Modality Integration, 4-Stage Reasoning, and Self-Refine—to 100 Hard samples.¹¹

¹¹**Modality Integration:** convert all forms of knowledge into text representations, employing chart captioning and table linearization (Yu et al., 2023b; Luo et al., 2023; Zhi Lim et al., 2024). **4-Stage Reasoning:** derives insights from Section 5.2

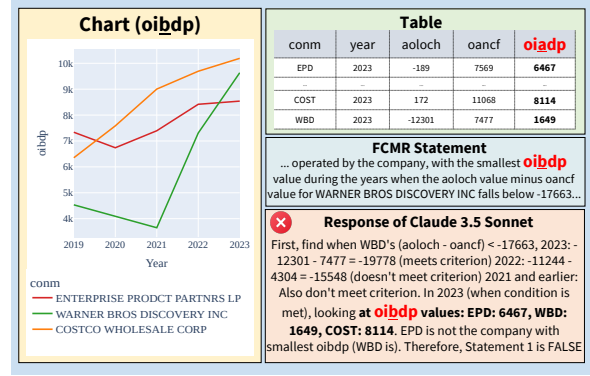


Figure 7: A case study on Claude’s errors: The condition, “the years when the aoloch value minus the oancf value for WARNER BROS DISCOVERY INC falls below -17663,” is satisfied for 2023. This task thus requires finding a company with a small oibdp value for 2023. Although the oibdp value is in the chart, not the table, Claude ignores the chart and incorrectly substitutes the oia dp value, likely due to the similar column names.

Method (Claude 3.5 Sonnet; tested on 100 Hard cases)	Acc. (%)
Zero-Shot COT (Baseline)	32
Modality Integration	39
Modality Integration + 4-Stage Reasoning	42
Modality Integration + 4-Stage Reasoning + Self-Refine	46

Table 8: Preliminary task optimization results.

The results confirm each method’s contribution to improved accuracy, verifying their effectiveness. This method also mitigates failure rates in both the Information Retrieval and Reasoning stages mentioned in Section 5.2 (see Figure 9). However, considering that the combination of all three methods saturates at 46% accuracy, future work is encouraged to develop dedicated approaches for FCMR.

7 Conclusion

We introduce FCMR, a new benchmark designed to evaluate the cross-modal multi-hop reasoning ability of MLLMs. We evaluate the performance of state-of-the-art MLLMs, revealing that current models struggle with reasoning across different modalities. In future work, we plan to develop optimized methods to enhance performance based on the observations and analyses from this study.

and explicitly mentions four reasoning stages in the prompt. **Self-Refine:** instructs the model to iteratively refine its own answer (Madaan et al., 2023). The prompts for each method are shown in Figure 25, Figure 26, and Figure 27, respectively.

Limitations

We present several points that can serve as the foundation for improving this work and initiating future research.

Potential for Extension to Other Domains

While we have conducted extensive experiments and analyses in the financial domain using FCMR, the proposed dataset generation framework, CM-RGen, has the potential to extend beyond the financial and material science domains, enabling the creation of datasets in fields such as law, biology, medicine, and electrical engineering. Future work can consider performing comprehensive performance evaluations of various models across these domains.

Reliance on Manual Analysis Part of the analysis in this work is based on the manual inspection of in-house researchers. While this was inevitable to guarantee a high-quality, in-depth investigation, future work may concentrate more on automated analysis. We also emphasize that, despite its cost, manual analysis is worth investigating, as demonstrated in Section 5.2 and Section 6.3.

Acknowledgments

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2025-00558151). This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) under the artificial intelligence semiconductor support program to nurture the best talents (IITP-(2025)-RS-2023-00253914) grant funded by the Korea government(MSIT). This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2020-II201373, Artificial Intelligence Graduate School Program(Hanyang University)).

References

- Amirhossein Abaskohi, Spandana Gella, Giuseppe Carenini, and Issam H Laradji. 2024. Fm2ds: Few-shot multimodal multihop data synthesis with knowledge distillation for question answering. *arXiv preprint arXiv:2412.07030*.
- Mubashara Akhtar, Chenxi Pang, Andreea Marzoca, Yasemin Altun, and Julian Martin Eisenschlos. 2024.
- Tanq: An open domain dataset of table answered questions. *arXiv preprint arXiv:2405.07765*.
- Amazon. 2024. [Amazon nova premier: Technical report and model card](#). *Amazon Technical Reports*.
- Anthropic. 2024. Claude: Large language model by anthropic. Available online at <https://www.anthropic.com>. Accessed on 2024-10-15.
- Yingshan Chang, Mridu Narang, Hisami Suzuki, Guihong Cao, Jianfeng Gao, and Yonatan Bisk. 2022. Webqa: Multihop and multimodal qa. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16495–16504.
- Wenhu Chen, Ming-Wei Chang, Eva Schlinger, William Yang Wang, and William W. Cohen. 2021a. [Open question answering over tables and text](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Yang Wang. 2020. [HybridQA: A dataset of multi-hop question answering over tabular and textual data](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1026–1036, Online. Association for Computational Linguistics.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. 2021b. [FinQA: A dataset of numerical reasoning over financial data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3697–3711, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Theodore E Christensen, Karson E Fronk, Joshua A Lee, and Karen K Nelson. 2024. Data visualization in 10-k filings. *Journal of Accounting and Economics*, 77(2-3):101631.
- Deepmind. 2025. [Gemini 2.0 flash thinking](#).
- DeepMind. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Yifan Du, Zikang Liu, Yifan Li, Wayne Xin Zhao, Yuqi Huo, Bingning Wang, Weipeng Chen, Zheng Liu, Zhongyuan Wang, and Ji-Rong Wen. 2025. [Virgo: A preliminary exploration on reproducing o1-like mllm](#). *Preprint*, arXiv:2501.01904.
- Ziliang Gan, Yu Lu, Dong Zhang, Haohan Li, Che Liu, Jian Liu, Ji Liu, Haipang Wu, Chaoyou Fu, Zenglin Xu, Rongjunchen Zhang, and Yong Dai.

2024. [Mme-finance: A multimodal finance benchmark for expert-level understanding and reasoning](#). *Preprint*, arXiv:2411.03314.
- Darryl Hannan, Akshay Jain, and Mohit Bansal. 2020. Manymodalqa: Modality disambiguation and qa over diverse inputs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7879–7886.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. 2023. [Seed-bench: Benchmarking multimodal llms with generative comprehension](#). *Preprint*, arXiv:2307.16125.
- Yongqi Li, Wenjie Li, and Liqiang Nie. 2022. Mmcoqa: Conversational question answering over text, tables, and images. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 4220–4231.
- Fangyu Liu, Julian Eisenschlos, Francesco Piccinno, Syrine Krichene, Chenxi Pang, Kenton Lee, Mandar Joshi, Wenhui Chen, Nigel Collier, and Yasemin Altun. 2023. [DePlot: One-shot visual language reasoning by plot-to-table translation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10381–10399, Toronto, Canada. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26296–26306.
- Timothy Liu and De Wen Soh. 2022. [Towards better characterization of paraphrases](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8592–8601, Dublin, Ireland. Association for Computational Linguistics.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2024b. Mmbench: Is your multi-modal model an all-around player? In *Computer Vision – ECCV 2024*, pages 216–233, Cham. Springer Nature Switzerland.
- Haohao Luo, Ying Shen, and Yang Deng. 2023. [Unifying text, tables, and images for multimodal question answering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 935–9367, Singapore. Association for Computational Linguistics.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). *Preprint*, arXiv:2303.17651.
- Ahmed Masry and Mehrad Shahmohammadi. 2024. [ChartInstruct: Instruction tuning for chart comprehension and reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 10387–10409, Bangkok, Thailand. Association for Computational Linguistics.
- OpenAI. 2024a. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- OpenAI. 2024b. [Learning to reason with llms](#).
- Chaoxu Pang, Yixuan Cao, Chunhao Yang, and Ping Luo. 2024. [Uncovering limitations of large language models in information seeking from tables](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 1388–1409. Association for Computational Linguistics.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yeqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Pooyan Rahmanzadehgervi, Logan Bolton, Mohammad Reza Taesiri, and Anh Totti Nguyen. 2024. Vision language models are blind. In *Proceedings of the Asian Conference on Computer Vision*, pages 18–34.
- Hossein Rajabzadeh, Suyuchen Wang, Hyock Ju Kwon, and Bang Liu. 2023. [Multimodal multi-hop question answering through a conversation between tools and efficiently finetuned large language models](#). *Preprint*, arXiv:2309.08922.
- Revant Gangi Reddy, Xilin Rui, Manling Li, Xudong Lin, Haoyang Wen, Jaemin Cho, Lifu Huang, Mohit Bansal, Avirup Sil, Shih-Fu Chang, et al. 2022. Mumuqa: Multimedia multi-hop news question answering via cross-media knowledge extraction and grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11200–11208.
- Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hananeh Hajishirzi, and Jonathan Berant. 2021. [Multimodalqa: complex question answering over text, tables and images](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal

- Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Haoran Wang, Aman Rangapur, Xiong Xiao Xu, Yue-qing Liang, Haroon Gharwi, Carl Yang, and Kai Shu. 2024. Piecing it all together: Verifying multi-hop multimodal claims. *arXiv preprint arXiv:2411.09547*.
- xAI. 2024. [Bringing grok to everyone](#).
- Qianqian Xie, Dong Li, Mengxi Xiao, Zihao Jiang, Ruoyu Xiang, Xiao Zhang, Zhengyu Chen, Yueru He, Weiguang Han, Yuzhe Yang, Shunian Chen, Yifei Zhang, Lihang Shen, Daniel Kim, Zhiwei Liu, Zhenheng Luo, Yangyang Yu, Yupeng Cao, Zhiyang Deng, Zhiyuan Yao, Haochang Li, Duanyu Feng, Yongfu Dai, Vijaya Sai Somasundaram, Peng Lu, Yilun Zhao, Yitao Long, Guojun Xiong, Kaleb Smith, Honghai Yu, Yanzhao Lai, Min Peng, Jianyun Nie, Jordan W. Suchow, Xiao-Yang Liu, Benyou Wang, Alejandro Lopez-Lira, Jimin Huang, and Sophia Ananiadou. 2024. [Open-finllms: Open multimodal large language models for financial applications](#). *Preprint*, arXiv:2408.11878.
- An Yang and Baosong Yang. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Huanjin Yao, Jiaxing Huang, Wenhao Wu, Jingyi Zhang, Yibo Wang, Shunyu Liu, Yingjie Wang, Yuxin Song, Haocheng Feng, Li Shen, and Dacheng Tao. 2024. [Mulberry: Empowering mllm with o1-like reasoning and reflection via collective monte carlo tree search](#). *Preprint*, arXiv:2412.18319.
- Yuan Yao and Tianyu Yu. 2024. [Minicpm-v: A gpt-4v level mllm on your phone](#). *Preprint*, arXiv:2408.01800.
- Kaining Ying, Fanqing Meng, Jin Wang, Zhiqian Li, Han Lin, Yue Yang, Hao Zhang, Wenbo Zhang, Yuqi Lin, Shuo Liu, Jiayi Lei, Quanfeng Lu, Runjian Chen, Peng Xu, Renrui Zhang, Haozhe Zhang, Peng Gao, Yali Wang, Yu Qiao, Ping Luo, Kaipeng Zhang, and Wenqi Shao. 2024. [MMT-bench: A comprehensive multimodal benchmark for evaluating large vision-language models towards multitask AGI](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 57116–57198. PMLR.
- Bowen Yu, Cheng Fu, Haiyang Yu, Fei Huang, and Yongbin Li. 2023a. [Unified language representation for question answering over text, tables, and images](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4756–4765, Toronto, Canada. Association for Computational Linguistics.
- Bowen Yu, Cheng Fu, Haiyang Yu, Fei Huang, and Yongbin Li. 2023b. [Unified language representation for question answering over text, tables, and images](#). *Preprint*, arXiv:2306.16762.
- Xiang Yue, Yuansheng Ni, Tianyu Zheng, Kai Zhang, Ruqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhao Chen. 2024. [Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9556–9567.
- Qing Zhang, Haocheng Lv, Jie Liu, Zhiyun Chen, Jianyong Duan, Hao Wang, Li He, and Mingying Xu. 2024a. An entailment tree generation approach for multimodal multi-hop question answering with mixture-of-experts and iterative feedback mechanism. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 4814–4822.
- Wenxuan Zhang, Sharifah Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2024b. M3exam: a multilingual, multimodal, multilevel benchmark for examining large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23, Red Hook, NY, USA. Curran Associates Inc.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. [PAWS: Paraphrase adversaries from word scrambling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bowen Zhao, Tianhao Cheng, Yuejie Zhang, Ying Cheng, Rui Feng, and Xiaobo Zhang. 2024. Ct2c-qa: Multimodal question answering over chinese text, table and chart. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 3897–3906.
- Qi Zhi Lim, Chin Poo Lee, Kian Ming Lim, and Ahmad Kamsani Samingam. 2024. [Unirag: Unification, retrieval, and generation for multimodal question answering with pre-trained language models](#). *IEEE Access*, 12:71505–71519.
- Fengbin Zhu, Wenqiang Lei, Yucheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. [TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, pages 3277–3287, Online. Association for Computational Linguistics.

A Details of Material Cross-Modal Multi-Hop Reasoning (Material-CMR)

The proposed CMRGen framework is easily adaptable to various domains, and as an example, we applied it to the field of Material Science to create the Material Cross-Modal Multi-Hop Reasoning (Material-CMR) dataset. We construct a Table Source containing material properties and a Text Source describing the materials through The Materials Project. The Materials Project¹² is an initiative that accelerates materials discovery by providing researchers with computational data and tools to predict material properties, enabling more targeted and efficient experimental research.

Specifically, we transform the entity corresponding to companies in FCMR into materials. Following the same steps proposed in CMRGen, we build Table Sources and Text Sources for Material-CMR. The Table Sources include columns for material properties such as band gap, density, volume, and formation energy per atom, while the Text Sources provide descriptions of the material’s crystal structure, structural characteristics, and properties. Using these Text Sources and Table Sources, we create datasets combining text, table, and chart input data. Subsequently, we generate single-modal one-hop, cross-modal two-hop, and cross-modal three-hop statements through GPT-4o-mini and templates, categorizing them by difficulty level. Also, A two-stage paraphrasing process is employed to maximize diversity. An example of the dataset instance is shown in Figure 8.

B Details on the Procedure of CMRGen

This section provides a comprehensive explanation of the dataset generation process for FCMR, detailing the sources and preprocessing steps for input table, text, and chart modalities, as well as the construction of table sources, text sources, and distractors.

B.1 Table Source

WRDS Compustat: Annual Simplified Financial Statements Considering the practicality of cross-modal multi-hop reasoning, we utilize *annual financial statements*, an essential element in real corporate analysis, as the table data source. Wharton Research Data Services (WRDS) Com-

pustat¹³ provides various financial data of publicly traded companies in North America. Among them, we use the **Annual Simplified Financial Statement**, which includes key financial columns such as Company Name, Ticker Code, Year, Net Sales, and Total Assets, spanning multiple years for each company. The Annual Simplified Financial Statement consists of 80 columns, which are broadly classified into four categories: Identifying Information, Balance Sheet Variables, Income Statement Variables, and Statement of Cash Flows Variables. The components of each category are presented in Figure 28. This Annual Simplified Financial Statement will later be used to construct the table and chart modalities.

Preprocessing The raw dataset contains a total of 80 columns. We standardize the unique symbol IDs to ticker codes and perform preprocessing to remove two columns that are not in millions of units to unify the units by column, leaving a total of 70 columns. We also use data from the most recent five years, 2019 to 2023.

B.2 Text Source

SEC EDGAR: 10-K Report To construct texts that are closely related to the Annual Simplified Financial Statement of company entities, we focus on corporate financial reports. Companies listed on the U.S. stock market are required to periodically provide financial reports to the U.S. Securities and Exchange Commission (SEC), and these reports can be publicly accessed through the Electronic Data Gathering, Analysis, and Retrieval System (EDGAR).¹⁴ We use the annual disclosure report, the 10-K report, of companies to match the Annual Simplified Financial Statement. This 10-K report differs from the summary-style annual reports typically used in datasets such as FinQA (Chen et al., 2021b) and TAT-QA (Zhu et al., 2021), as it provides more in-depth financial data and disclosures. All companies’ 10-K reports are composed of a common table of contents format.

Each 10-K report includes several key items that are vital for corporate analysis. For instance, Item 1 provides a description of the company’s business model, its products or services, and its primary markets. Item 7, often referred to as the Management’s Discussion and Analysis (MD&A), allows

¹²<https://next-gen.materialsproject.org/>

¹³<https://wrds-www.wharton.upenn.edu/pages/grid-items/compustat-global-wrds-basics/>

¹⁴<https://www.sec.gov/search-filings>

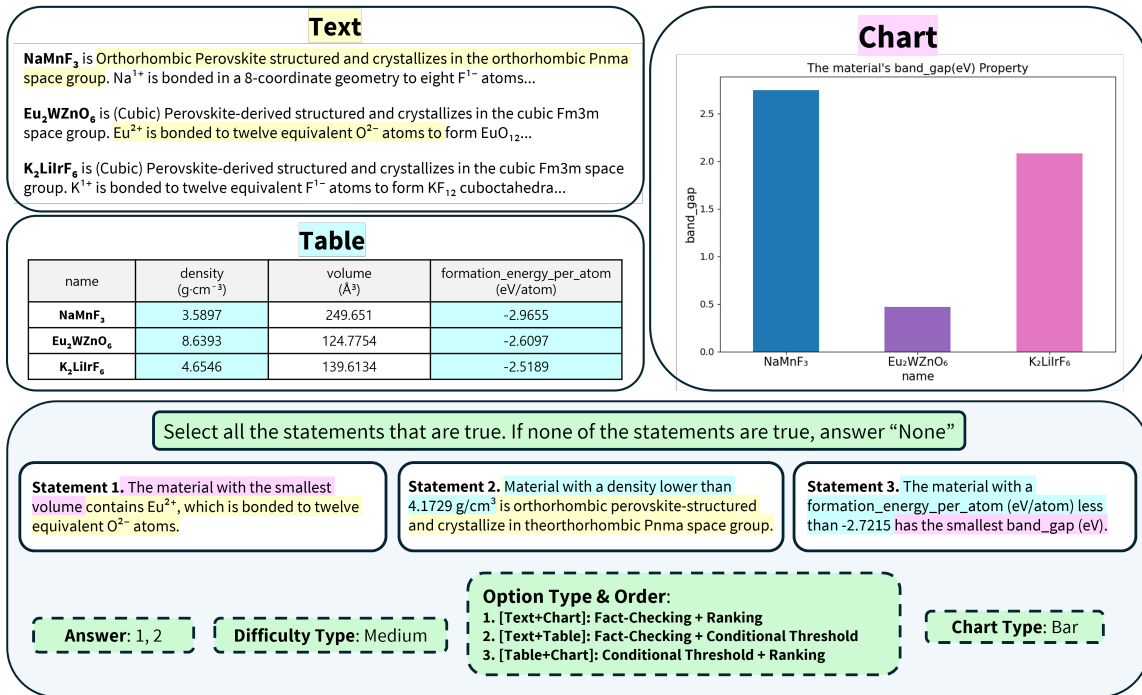


Figure 8: An example from Material-CMR.

company executives to discuss operational results, providing insight into trends, risks, and strategies. In addition, Item 7A covers quantitative and qualitative disclosures regarding market risks, while Item 8 presents the audited financial statements, offering a transparent view of the company’s financial health. These items, along with other sections, make the 10-K an essential document for evaluating a company’s long-term viability and strategy.

Among the various items of the 10-K report, we use ITEM 1 (Business), ITEM 2 (Legal Proceedings), ITEM 7 (Management’s Discussion and Analysis of Financial Condition and Results of Operations), ITEM 7A (Quantitative and Qualitative Disclosures about Market Risk), and ITEM 8 (Financial Statements and Supplementary Data), which are most commonly used in actual corporate analysis. Further details can be found in the document provided by the SEC¹⁵.

This 10-K report data will later be used to construct the input text.

Preprocessing To align with the Annual Simplified Financial Statement data, we filter companies where both the Annual Simplified Financial Statements and 10-K annual reports exist, ensuring all formats of reports for the most recent five years are fully present. Among them, we select the top 101

companies based on Net Sales in 2023.

B.3 Table Source Construction

A table source serves as an intermediate bridge connecting the text, table, and chart modalities and is used as a base anchor for creating multi-hop statements. After sampling three companies from the Annual Simplified Financial Statement data of the 101 companies, we construct the table source by randomly sampling seven financial columns excluding company name, ticker code, and year. One column, used for chart generation, is chosen to avoid NaN values. The final generated table source consists of the Annual Simplified Financial Statement data of three companies, each having five years, and is composed of ten columns.

B.4 Text Source Construction

The 10-K reports obtained through SEC EDGAR are too lengthy to use entirely at once as input text. Therefore, we divide the 10-K reports of each company into chunks of three consecutive paragraphs. Later, these chunks will be used as the input text.

B.5 Input Data Construction: Text, Tables, and Charts

A single data instance contains a total of three companies. From the table source with three companies, we select one column without NaN values

¹⁵<https://www.sec.gov/files/reada10k.pdf>

as the chart column and convert the table source excluding the chart column into the input table and the chart column into the input chart. To preserve the structural information of the table, the input table is constructed in JSON format, and to ensure data diversity, the input chart uses three different libraries and four chart types (line, bar, scatter, pie) commonly used in financial domains. The input text corresponds to the text sources of the three companies. All of these processes are automated through a Python script.

B.6 Input Chart Code Generation

The input chart in FCMR consists of four types: Line, Bar, Scatter, and Pie, generated using visualization libraries such as matplotlib, seaborn, and plotly. To enhance chart diversity and mitigate data bias, 16 font types, including ['Arial', 'Verdana', 'Times New Roman', 'Courier New', 'Georgia', 'Comic Sans MS', 'Tahoma', 'Cambria', 'Microsoft YaHei', 'Nirmala UI', 'Calibri', 'Consolas', 'Segoe UI', 'Garamond', 'Century Schoolbook', 'Book Antiqua'], were applied to text within the charts. The font size for titles, labels, legends, and ticks was randomly selected within predefined minimum and maximum thresholds. To ensure clear visual distinction, the color palette consisted of seven colors: ['#1f77b4', '#ff7f0e', '#2ca02c', '#d62728', '#9467bd', '#8c564b', '#e377c2']. The thickness of lines and bars was also randomly selected within predefined thresholds. To clearly visualize trends and rankings in charts, we introduced controlled variance in yearly data values, ensuring the design avoids cases where differentiation is visually ambiguous.

C Details of Data Quality Control

C.1 Paraphrasing Quality

To evaluate the quality of Lexical-Syntax 2-Stage Paraphrasing, we employed the Word Position Deviation (WPD) and Lexical Deviation (LD) metrics proposed in (Liu and Soh, 2022). WPD assesses the syntactic diversity of paraphrased sentences, while LD evaluates lexical diversity. For an objective comparison, as shown in Table 3, we benchmarked the WPD and LD metrics of FCMR against prominent paraphrasing datasets such as MRPC (Dolan and Brockett, 2005) and PAWS (Zhang et al., 2019), demonstrating the superior quality of our paraphrasing. Additionally, we validated semantic consistency using Claude 3.5 Sonnet to

filter out samples where the paraphrased sentences were flagged as semantically altered.

To ensure high quality and robustness in FCMR’s Hard-level instances, a specific human annotation and review process was employed. Human reviewers initially examined 10% of samples from all difficulty levels (Easy, Medium, and Hard). This initial review showed that Easy and Medium samples were of high quality, so no further extensive human annotation was deemed necessary for them. Due to their complexity, Hard-level questions underwent a more rigorous human review to ensure their quality. Through this focused review, 22 out of 714 Hard-level samples (approximately 3%) were identified for revision. These revisions aimed to improve clarity, correctness, or the intended reasoning complexity. Finally, two authors cross-checked each revision to ensure consensus and maintain high-quality standards.

C.2 Verification of Data Contamination

To ensure a fair comparison of data contamination between MMQA and FCMR under identical conditions, we evaluated instances requiring cross-modal three-hop reasoning from each dataset using the GPT-4o model under the following settings: (1) Random Selection, (2) Without Image Input, and (3) With Image Input. For MMQA, *Random Selection* in Table 2 involves randomly selecting a single word from the question, text, or table. In contrast, for FCMR, *Random Selection* in Table 4 involves randomly selecting one of the eight possible answers, ranging from none to (1, 2, 3).

C.3 Bias Mitigation Strategies

Due to the design requiring reasoning over three statements, there is a potential for bias to arise from specific factors. To minimize bias, we implemented several strategies. First, we ensured a balanced distribution of modality order types to prevent bias toward specific order configurations. Second, we adjusted the distribution of statement types to avoid overrepresentation of particular types. Third, we maintained an even distribution across the eight answer types to reduce bias toward any specific answer type. The distributions of answer type, statement type, and library type across all difficulty levels are visualized in Figure 21.

D Case Study Examples

D.1 Trend Assessment Error

As shown in Figure 14, Claude struggles to identify increasing trends. This difficulty is particularly pronounced when interpreting cumulative bar charts or charts with ranges that include negative values, where the success rate of interpretation is significantly lower.

D.2 Ranking/Ordering Mistake

Figure 15 illustrates a case where the Claude model fails to accurately determine the ranking for a specific year from a chart. While the model performs better in identifying rankings compared to recognizing increasing or decreasing trends, its success rate remains significantly lower when interpreting cumulative bar charts or charts with ranges that include negative values.

D.3 Condition Satisfaction & Selection Error

The model sometimes fails to correctly identify a company or element that meets given conditions, or asserting that no such entity exists when one does. An example is in Figure 16.

D.4 Data/Value Interpretation Error

The model occasionally fails in calculations involving addition when negative numbers are included or when the number of terms exceeds three. Additionally, there are instances where it fails to correctly compare the magnitude of numbers. Figure 17 illustrates one such case. Considering that addition and magnitude comparison are simple operations for humans, this highlights the need for improvement in the arithmetic reasoning capabilities of MLLMs.

D.5 Company/Entity Confusion

Errors in this category involve mixing up one company or entity with another. Even when companies are distinguished by unique colors, labels, or legends, the model may incorrectly assign data from one company to a different one, thus undermining the validity of its reasoning and final answers. An example is in Figure 18.

D.6 Information Extraction & Reasoning Failure

There are the cases incorrectly extracting facts, misunderstanding textual information, drawing un-

Model	Accuracy (%)
Claude 3.5 Zero-Shot COT Baseline	32
Mulberry-LLaVA-8B (Yao et al., 2024)	12
Virgo 72B (Du et al., 2025)	19
Gemini 2.0 Flash Thinking (Deepmind, 2025)	39
o1 (OpenAI, 2024b)	43
nova-pro-v1 (Amazon, 2024)	22
grok-2-vision (xAI, 2024)	19
qwen2.5-vl-72b-instruct (Qwen et al., 2025)	19
gemma-3-27b-it (DeepMind, 2025)	19

Table 9: Performance of various advanced Multimodal Large Language Models (MLLMs), including o1-like reasoning models and other recent models, on 100 randomly selected FCMR Hard-level samples. Evaluation was limited to this 100-sample subset due to the operational costs associated with these models.

justified conclusions, or logical missteps after gathering correct details. An example is in Figure 19.

D.7 Problem Understanding & Condition Ignoring

Claude sometimes makes incorrect judgments by considering only a subset of the required conditions. This issue is particularly prominent in Hard-level tasks that require deep reasoning. An example of this case is in Figure 20.

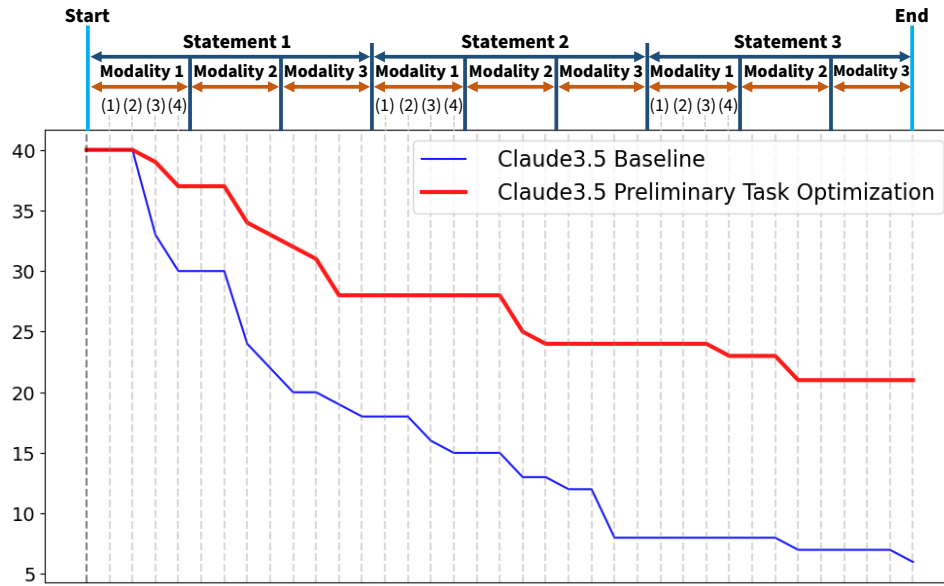


Figure 9: Fine-grained Stage-based Analysis of Preliminary Task Optimization results applying Modality Integration, 4-Stage Reasoning, and Self-Refine.

MMQA Data Contamination Example

Question: Which team with scoring leaders in the 2004-05 AHL season has a lower PIM: The team pk subban's brother plays for or the team with a logo featuring a skull?

Text

Pernell-Karl ... professional ice hockey defenceman for the New Jersey Devils of the National Hockey League (NHL).... **Malcolm-Jamaal Justin Subban** (born December 21, 1993) is a Canadian ice hockey goaltender. He currently plays in the American Hockey League with the **Providence Bruins** as a prospect of.... Hockey League with the Belleville Bulls. **His older brother, P. K. Subban** is a defenceman with the Nashville Predators.

Table

Player	Team	GP	G	A	Pts	PIM
Jason Spezza	Binghamton Senators	80	32	85	117	50
Mike Cammalleri	Manchester Monarchs	79	46	63	109	60
David Ling	St. John's Maple Leafs	80	28	60	88	152
Kyle Wellwood	St. John's Maple Leafs	80	38	49	87	20
Simon Gamache	Milwaukee Admirals	80	29	57	86	93
Peter Sarno	Manitoba Moose	80	16	66	82	53
Chris Taylor	Rochester Americans	79	21	58	79	50
Andy Hilbert	Providence Bruins	79	37	42	79	83
Denis Hamel	Binghamton Senators	80	39	39	78	75
Eric Staal	Lowell Lock Monsters	77	26	51	77	88

Image

Answer of GPT-4o: Without Image Input

To determine which team has a lower PIM (Penalty In Minutes) between the team Malcolm-Jamaal Justin Subban's brother plays for and the team with a logo featuring a skull, we need to look at the details provided: 1. **Malcolm-Jamaal Justin Subban** is noted to play for the **Providence Bruins** in the American Hockey League (AHL) as a prospect for the Boston Bruins. 2. Looking at the AHL scoring leaders' table for the 2004-05 season: - **Providence Bruins** have a PIM of 83.0 (associated with Andy Hilbert). 3. **For the team with a logo featuring a skull, it would be the Milwaukee Admirals, as they are known for such a logo. - Milwaukee Admirals**...

Figure 10: Example of data contamination in MMQA. Despite the requirement to infer the pink-highlighted information, "team with a logo featuring a skull", using image data, GPT-4o deduced the answer, "Milwaukee Admirals", without being provided with the image. This suggests the model leveraged internal knowledge, rather than cross-modal reasoning, highlighting a limitation in evaluating the cross-modal multi-hop reasoning capabilities of MMQA.

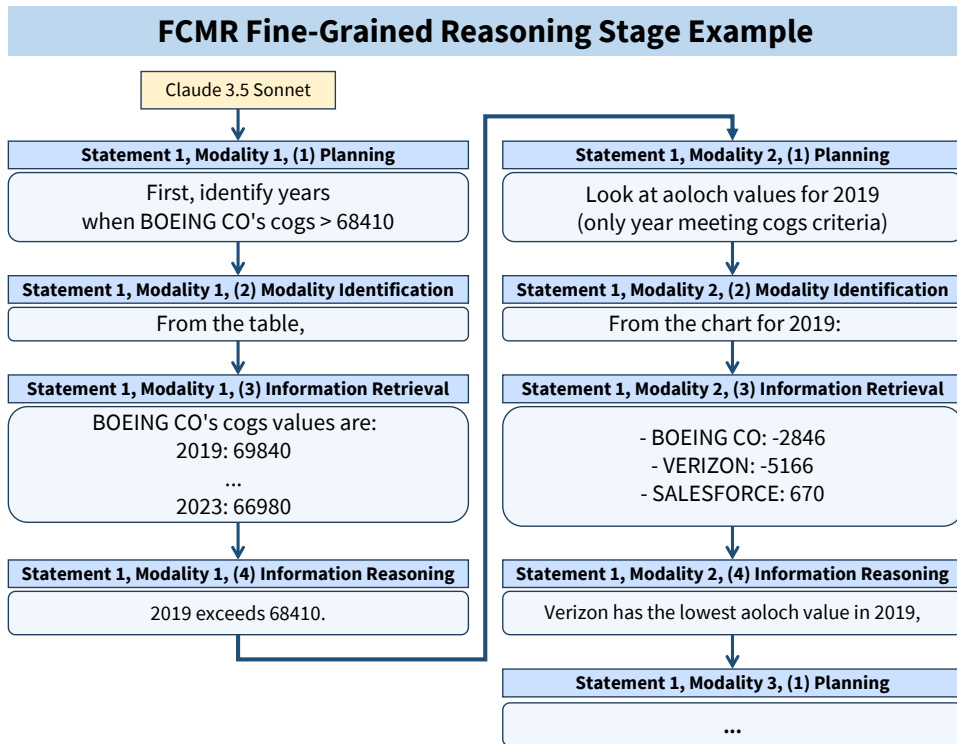


Figure 11: An example of decomposing the reasoning process of the Claude 3.5 Sonnet's response into fine-grained, stage-based steps.

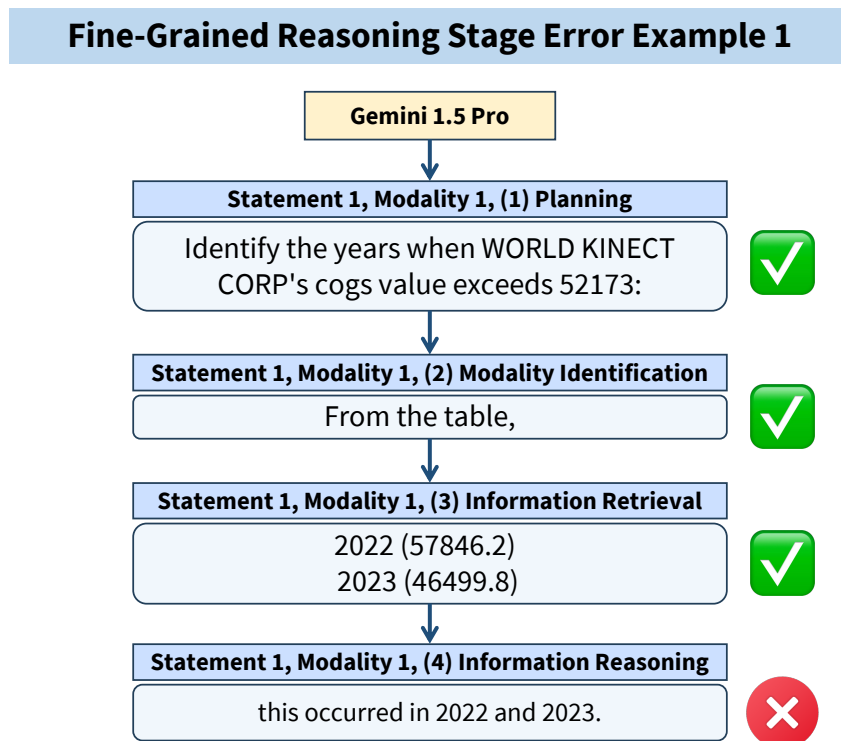


Figure 12: An example where the Gemini 1.5 Pro succeeds in (3) Information Retrieval but fails in (4) Information Reasoning. While the model successfully extracts the information that the cogs value is 57846.2 for 2022 and 46499.8 for 2023 from the table, it incorrectly reasons that the cogs values for both 2022 and 2023 are greater than 52173, resulting in a failure.

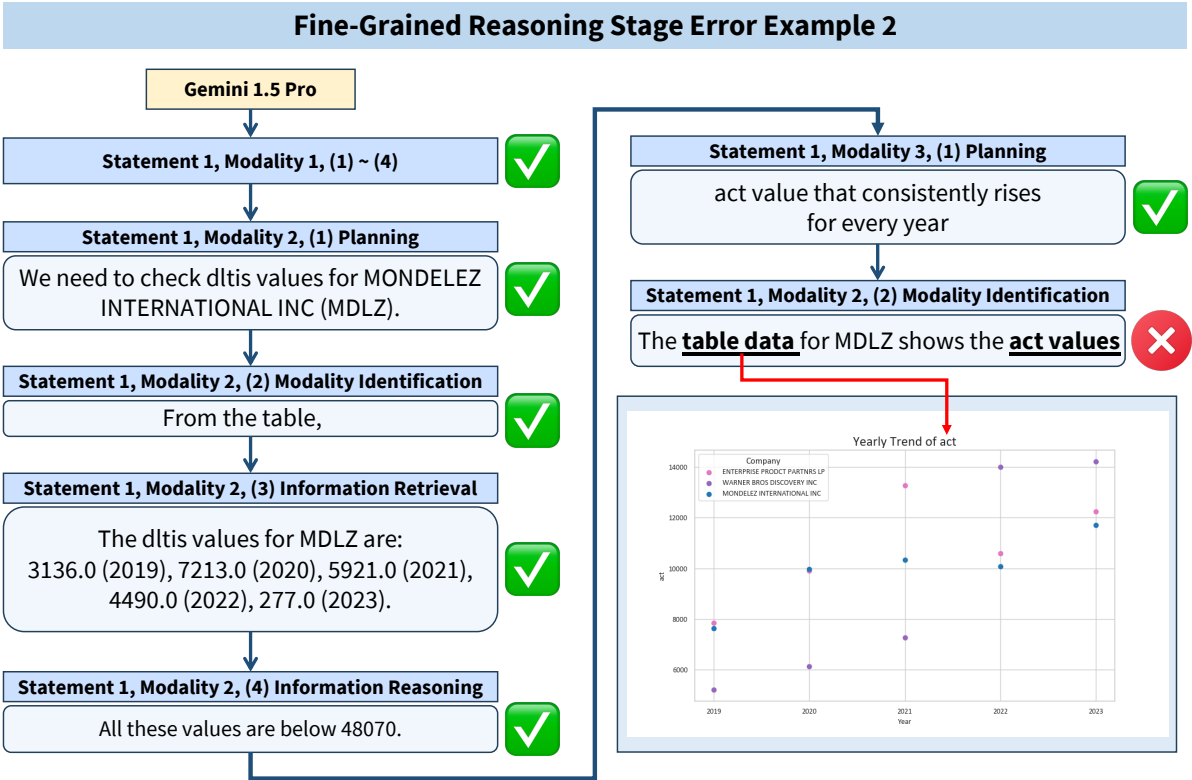


Figure 13: An example where the Gemini 1.5 Pro succeeds in (1) Planning but fails in (2) Modality Identification. While the model successfully plans that the act value is needed, it fails by identifying the modality as Table instead of Chart, where the act value is actually presented.

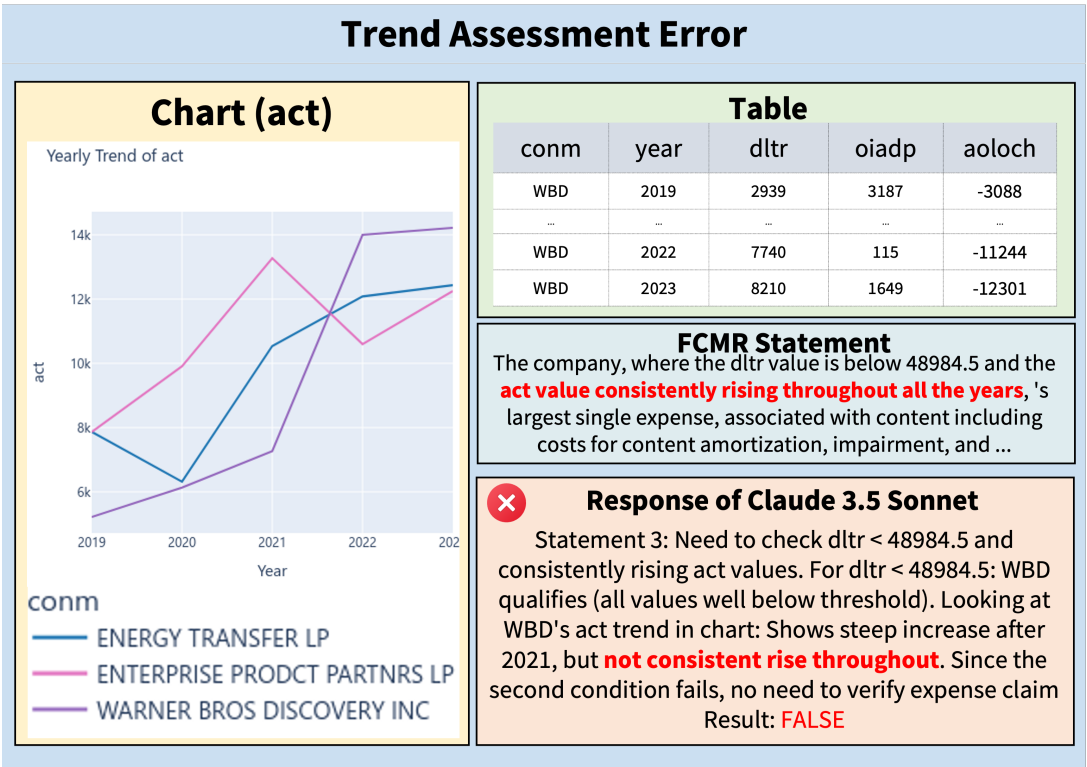


Figure 14: In the chart, WBD shows a consistently increasing trend in act values across all years, but Claude 3.5 Sonnet fails to recognize this correctly, resulting in an error.

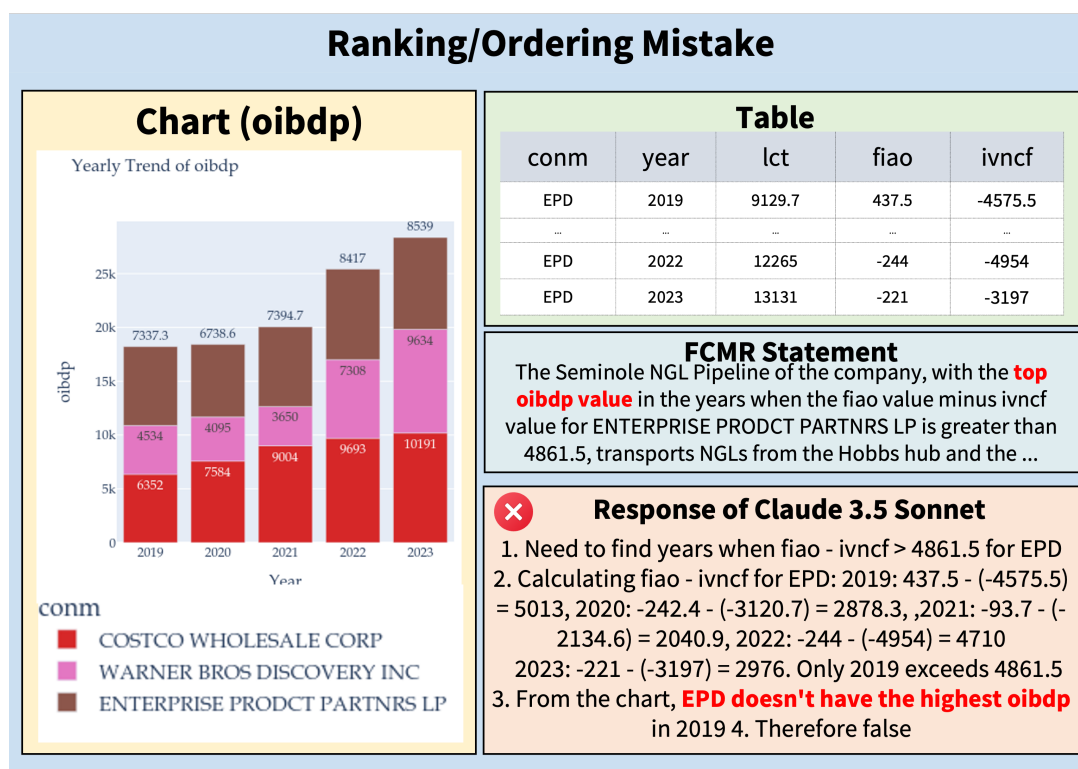


Figure 15: An example where the Claude model fails to interpret a cumulative graph: Although the company with the highest OIBDP value in 2019 is indeed EPD (Enterprise Product Partners LP), Claude incorrectly states otherwise.

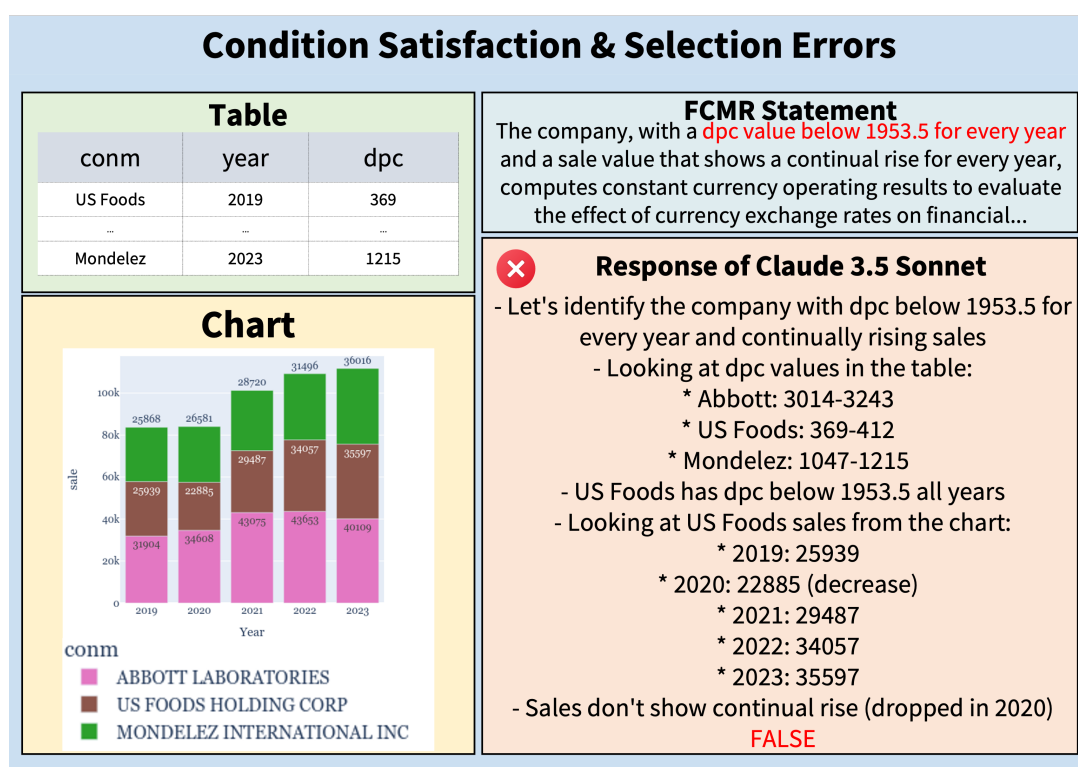


Figure 16: The condition that the dpc value is less than 1953.5 for all years is satisfied by both US Foods and Mondelez, but Claude 3.5 Sonnet recognizes only US Foods and fails to consider Mondelez, resulting in an error.

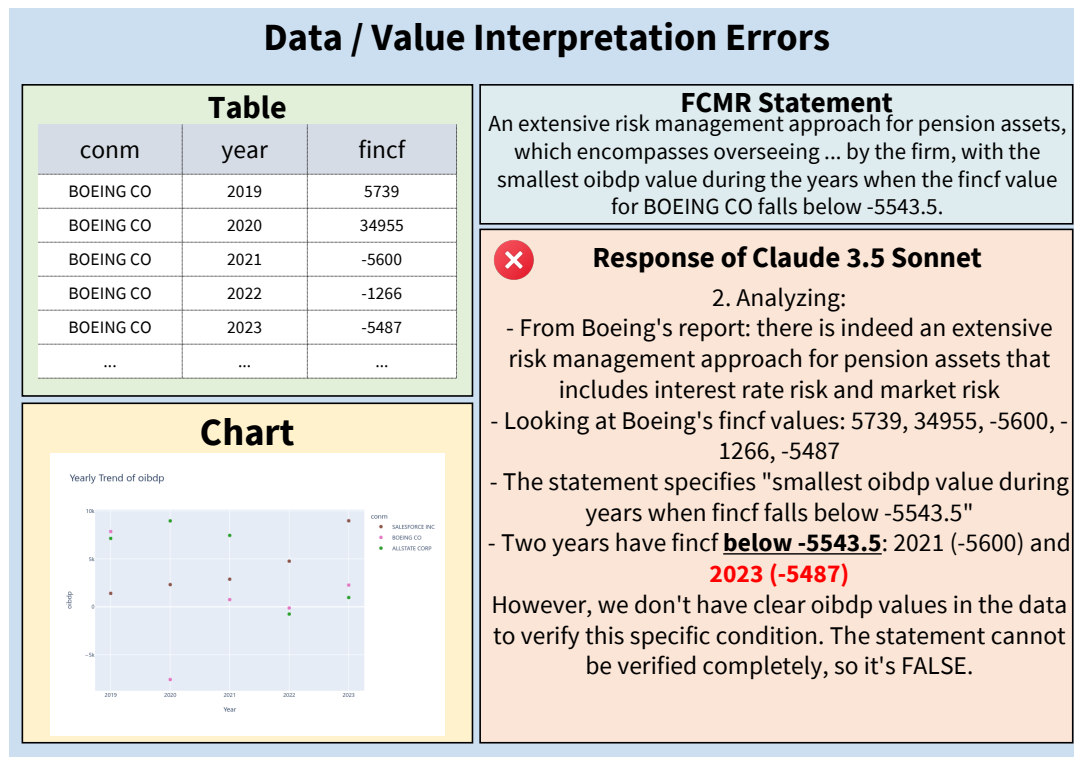


Figure 17: -5487 is not smaller than -5543.5, but Claude failed to correctly compare the two numbers and consequently reached the wrong conclusion.




Figure 18: Although charts and tables were successfully interpreted to determine that Salesforce had the highest act value in 2019, the company was ultimately confused in the final step by Claude, resulting in an incorrect conclusion.

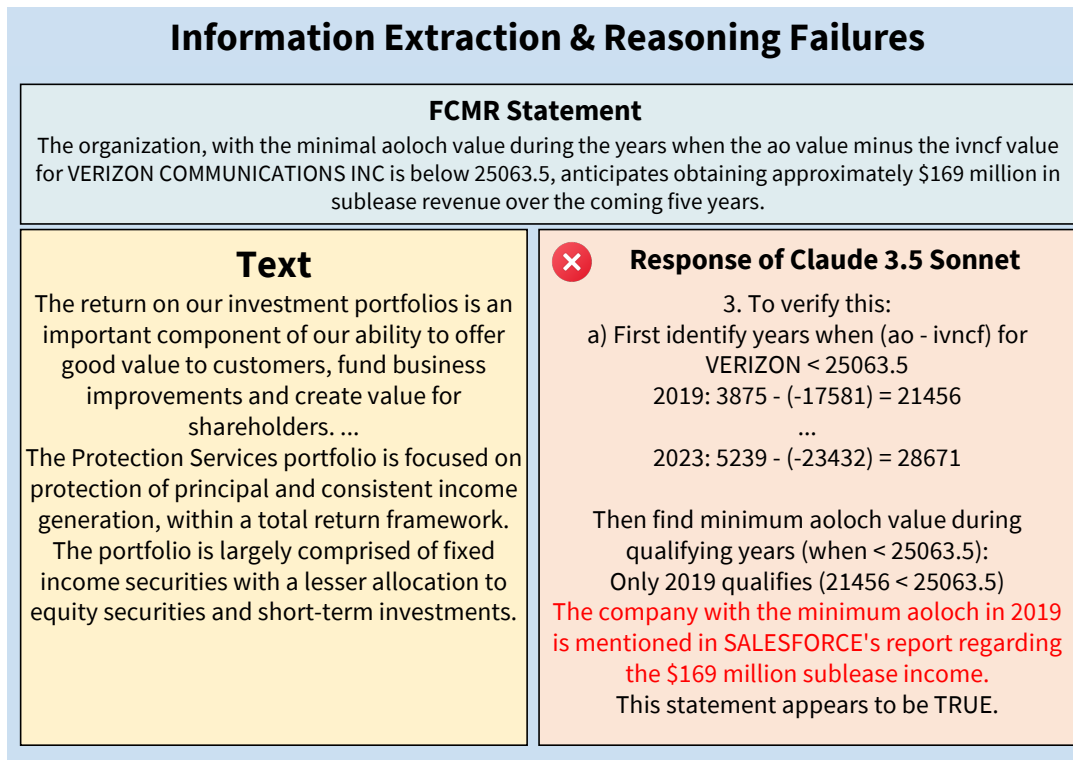


Figure 19: To compare the aoloch value, the chart should have been considered. Instead, having relied solely on the text data, Claude led to an incorrect conclusion.

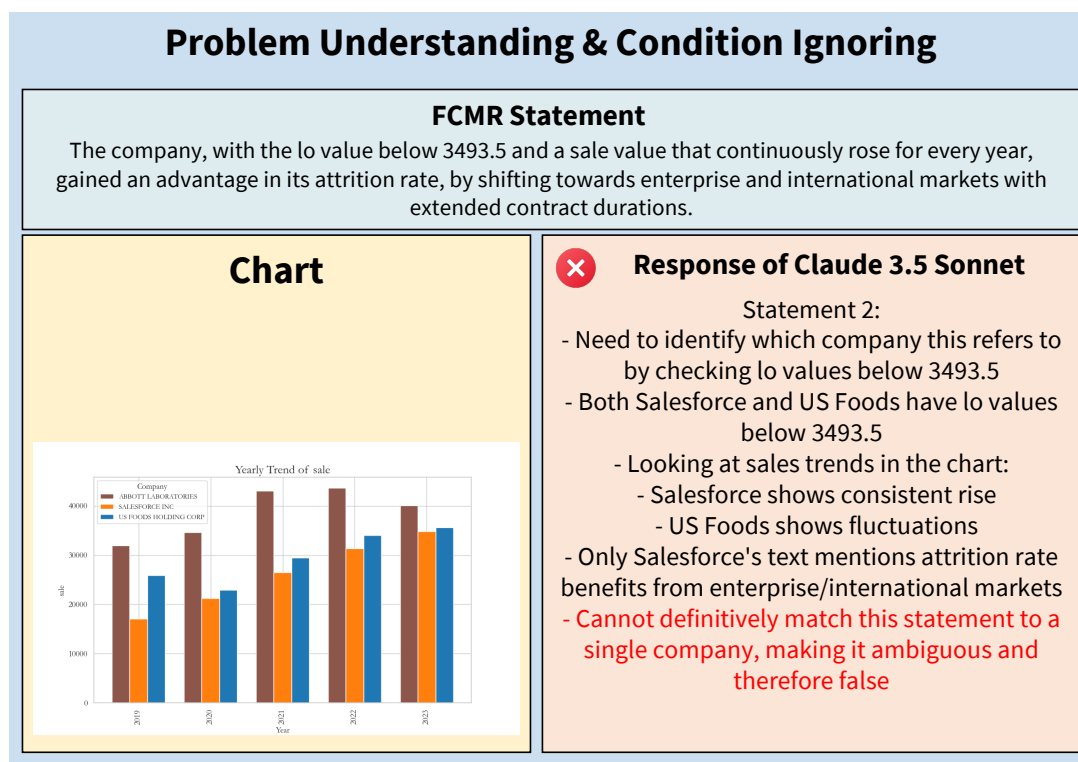


Figure 20: Although Claude found that Salesforce satisfied all the conditions, it misunderstood the statement and consequently made an incorrect judgment.




Figure 21: Pie charts for answer types, statement types, and library usage categorized by difficulty levels.




Figure 22: Distribution of inference failures for each MLLM across the four reasoning stages (Planning, Modality Identification, Information Retrieval, and Information Reasoning), based on manual analysis of 40 Hard-level samples.

Statement Types	Example
FC	In May 2020, WORLD KINECT CORP modified and refreshed its asset-backed debt financing facility.
CT	The firm COSTCO WHOLESALE CORP. disclosed that its cogs values surpassed 92,765 during the year 2021.
AR	The company where the 2023 txt value minus the 2022 txt value equals 359.0 is WORLD KINECT CORP.
TR	Over the period from 2020 through 2023, SALESFORCE INC. consistently experienced an increase in its xsga values.
RK	During 2019, US FOODS HOLDING CORP. possessed the lowest txt value compared to other companies.
FC+CT	In 2023, the company whose xint values are less than 849.5 owns a terminal facility located at Fort Mifflin, which includes two docks for ships and has a total storage capacity of approximately 570 MBbls.
FC+AR	The company where the 2020 spi value minus the 2022 spi value equals 1036.0 recorded \$39 million and \$40 million in prior service credit amortization in 2018 and 2017, respectively.
FC+TR	The firm that showed a steady increase in dpc values between 2019 and 2023 must comply with the detailed regulations set by the Department of Transportation (DOT) regarding its pipeline infrastructure.
FC+RK	In 2022, the company, which reports fincf values greater than -2009, spreads out the amortization of its capitalized costs tied to new revenue contracts across four years.
CT+TR	Over the years, MONDELEZ INTERNATIONAL INC. consistently reports a fincf value exceeding -19575.5, while the ceq value demonstrates a continuous increase.
CT+RK	When the seq figure for WARNER BROS DISCOVERY INC falls below 10177.5, VERIZON COMMUNICATIONS INC records the lowest cogs value.
AR+TR	Throughout all periods, ENTERPRISE PRODCOT PARTNRS LP is the company in which the cumulative nopi values surpass 1553.35, while the ao values have persistently increased.
AR+RK	During the years when lt value minus ibc value for UNITED PARCEL SERVICE INC is greater than 55265, the company with the lowest sale value is UNITED PARCEL SERVICE INC.
FC+CT+TR	Professional services are provided by the organization, which has the nopi value below 2796 for all years and the aoloch value that consistently declines for all years, to help clients with digital transformations using Salesforce solutions.
FC+CT+RK	The firm, with the highest act value during the years when ENERGY TRANSFER LP's intan value dips below 8059, acquired a controlling interest in USAC through a \$250 million cash transaction.
FC+AR+TR	The organization, with the cumulative sum of nopi values below 2840 and continuously increasing act values for every year, has provided put rights to certain consolidated subsidiaries. These put rights are omitted from the contractual obligations table due to unpredictability in payment.
FC+AR+RK	The business, with the minimal aoloch value in the years when the ivao value minus the ao value for VERIZON COMMUNICATIONS INC exceeds 9312.5, provides expert services to support customers in executing digital transformations leveraging Salesforce solutions.

Table 10: Examples for each statement type. FC refers to Fact-Checking, CT refers to Conditional Threshold, AR refers to Arithmetic, TR refers to Trend, and RK refers to Ranking.

Statement Types	Template
CT	The company with (column) values (greater than, less than) (threshold) in (Year) is company.
AR	The company where the (Year1) (column) value (plus, minus) the (Year2) (column) value equals (results) is (company).
TR	The company that showed a continuously (increasing, decreasing) trend in (column) values from (Year1) to (Year2) is (company).
RK	The company with the (highest, lowest) (column) value in (Year) is (company).
CT+TR	For all years, the company with the (column1) value (greater than, less than) (threshold) and the (column2) value continuously (increased, decreased) is (company).
CT+RK	During the years when the (column1) value for (company) is (greater than, less than) (threshold), the company with the (highest, lowest) (column2) value is (company).
AR+TR	For all years, the company with the cumulative sum of (column1) values (greater than, less than) (threshold) and the (column2) values continuously (increased, decreased) is (company).
AR+RK	During the years when (column1) value (plus, minus) (column2) value for (company1) is (greater than, less than) (threshold), the company with the (highest, lowest) (column3) value is (company2).

Table 11: Base templates of statement types. FC refers to Fact-Checking, CT refers to Conditional Threshold, AR refers to Arithmetic, TR refers to Trend, and RK refers to Ranking. In the case of Statement Types that include FC, a new template is generated by combining them with other Statement Types and Facts, where they share a common company entity.

Difficulty	Modality Types	Statement Types
Easy	Text	Fact-Checking
	Table	Conditional Threshold
		Arithmetic
	Chart	Trend Ranking
Medium	Text + Table	Fact-Checking + Conditional Threshold
		Fact-Checking + Arithmetic
	Text + Chart	Fact-Checking + Trend
		Fact-Checking + Ranking
	Table + Chart	Conditional Threshold + Trend
		Conditional Threshold + Ranking
Hard		Arithmetic + Trend
		Arithmetic + Ranking
	Text + Table + Chart	Fact-Checking + Conditional Threshold + Trend
		Fact-Checking + Conditional Threshold + Ranking
		Fact-Checking + Arithmetic + Trend
		Fact-Checking + Arithmetic + Ranking

Table 12: Detailed Statement Types by Difficulty and Modality Types. For the Easy level, all three answer statements are single-modal one-hop, while for the Medium level, all three statements are cross-modal two-hop. At the Hard level, all three statements consist of cross-modal three-hop. For Easy level questions, while each of the three statements is single-modal one-hop, determining the correct overall answer for the instance (i.e., selecting all true statements) requires processing information from all three provided modalities (text, table, and chart), ensuring comprehensive multi-modal engagement. Specific examples can be found in Table 10.

You are provided with the following materials:

[Text Reports: Detailed excerpts from company reports of three companies]

{text}

[Table Data: A table containing financial data for the same three companies]

{table}

[Chart Image]

{chart}

Your Task:

Determine whether each of the following three statements is true or false based solely on the provided materials. For each statement:

1. {option1}

2. {option2}

3. {option3}

Provide a detailed reasoning process that references specific data or information from the text reports, table data, or chart images.

Do not use general knowledge or external information beyond what is provided in the materials.

If there is insufficient information to determine the truthfulness of a statement, or if the statement relies on information not present in the materials, consider it false.

Final Answer Format:

After your reasoning, provide the final answer by listing the numbers of the statements that are true.

For example: "Answer: 1 or 1,2 or 2,3".

If none of the statements are true, write: "Answer: None".

Figure 23: MLLMs zero-shot prompt.

You are provided with the following materials:

[Text Reports: Detailed excerpts from company reports of three companies]
{text}

[Table Data 1: A table containing financial data for the same three companies]
{table}

[Table Data 2: A table containing financial data for the same three companies]
{chart_to_table}

Your Task:

Determine whether each of the following three statements is true or false based solely on the provided materials. For each statement:

1. {option1}
2. {option2}
3. {option3}

Provide a detailed reasoning process that references specific data or from the text reports, table data.

Do not use general knowledge or external information beyond what is provided in the materials.

If there is insufficient information to determine the truthfulness of a statement, or if the statement relies on information not present in the materials, consider it false.

Final Answer Format:

After your reasoning, provide the final answer by listing the numbers of the statements that are true.

For example: "Answer: 1 or 1,2 or 2,3".

If none of the statements are true, write: "Answer: None".

Figure 24: (M)LLMs + Deplot zero-shot prompt.

<p>Task: Create a comprehensive, data-rich caption for the provided chart. Your caption must include every visible detail and piece of information on the chart, with no information lost. Follow the instructions below strictly and in order.</p> <p>1. Identify the Chart Type and Overall Context Clearly state the type of chart (e.g., bar chart, line chart, scatter plot, pie chart, etc.). Mention the general subject or theme (e.g., sales over time, population distribution, survey results, etc.). Indicate the time period, region, or other contextual details if visible on the chart.</p> <p>2. Describe the Axes or Segments Name the x-axis (or horizontal axis) label, its units, and range if present. Name the y-axis (or vertical axis) label, its units, and range if present. For other types of charts (e.g., pie charts), describe how each segment or slice is labeled and what each label means.</p> <p>3. Detail Each Data Point or Category List out all categories or groups shown on the chart. Specify the corresponding data values for each category or group. Use exact numerical values if visible, or approximate them if only visually estimated. Highlight any color-coding, pattern fills, or legends, explaining what each color or pattern represents.</p> <p>4. Identify Trends and Patterns Explain any apparent increases, decreases, spikes, or plateaus in the data. State how categories compare to each other (e.g., which is highest, which is lowest, any ties, etc.). Note any step changes, anomalies, or outliers, and specify the relevant data points or labels.</p>	<p>5. Describe Key Highlights, Anomalies, or Outliers Mention any unusual shapes or points that deviate significantly from the rest of the data. Clearly specify where these outliers occur, and provide potential context or interesting observations if given by the chart or accompanying data.</p> <p>6. Include Additional Chart Elements Discuss the chart title and subtitle. Mention any annotations, data labels, or callouts on the chart. Describe any footnotes, data sources, or disclaimers shown.</p> <p>7. Provide Summarizing Insights Give a concise interpretation of what the chart suggests or proves, based on visible data. Compare the highest and lowest points, overall trends, and any significant patterns. If the chart compares multiple variables, mention how they relate to each other.</p> <p>8. Use Clear, Structured Language Present the information in a logically organized manner, using complete sentences. Keep the language as factual and objective as possible. Avoid assuming causes or motivations behind the data unless explicitly stated on the chart.</p> <p>9. No Information Loss Ensure that every visible element on the chart is mentioned. If exact figures are present, include them. If certain details are not clear or partially visible, state that explicitly. Goal: After following these instructions, you will produce an extensive, thorough caption that leaves no data point, visual cue, or contextual note unmentioned.</p>
---	--

Figure 25: Modality Integration prompt.

<p>For each statement, carry out your reasoning using the following strategy:</p> <ol style="list-style-type: none"> (1) Planning: identifying the required values. (2) Modality Identification: recognizing which modality contains these values. (3) Information Retrieval: extracting relevant information from the identified modality. (4) Information Reasoning: reasoning over the extracted information under the given conditions.

Figure 26: 4-Stage Reasoning Strategy prompt.

<p>Reasoning Agent made the following reasoning about the problem:</p> <p>{Initial Output}</p> <p>You are the Refinement Agent tasked with refining Reasoning Agent's reasoning. If you find any errors in Reasoning Agent's reasoning for each statement, provide a corrected answer along with appropriate reasoning. If no errors are found, retain the original answer.</p>
--

Figure 27: Self-Refine prompt.

<p><u>Identifying Information</u></p> <p>CUSIP (cusip)Ticker Symbol (tic)</p> <p>CIK Number (cik)</p> <p>Company Name (conm)</p> <p>Fiscal Year-end Month (fyr)</p> <p>State/Province (state)</p> <p>Standard Industry Classification Code (sic)</p> <p>North American Industry Classification Code (naics)</p> <p><u>Balance Sheet Variables</u></p> <p>Cash and Short-Term Investments (che)</p> <p>Receivables - Total (rect)</p> <p>Inventories - Total (invt)</p> <p>Current Assets - Other - Total (aco)</p> <p>Current Assets - Total (act)</p> <p>Property, Plant and Equipment - Total (Net) (ppent)</p> <p>Depreciation, Depletion and Amortization (Accumulated) (dpact)</p> <p>Investment and Advances - Equity (ivaeq)</p> <p>Investment and Advances - Other (ivao)</p> <p>Intangible Assets - Total (intan)</p> <p>Assets - Other (ao)</p> <p>Assets - Total (at)</p> <p>Debt in Current Liabilities - Total (dlc)</p> <p>Accounts Payable - Trade (ap)</p> <p>Income Taxes Payable (txp)</p> <p>Current Liabilities - Other - Total (lco)</p> <p>Current Liabilities - Total (lct)</p> <p>Long-Term Debt - Total (dltt)</p> <p>Deferred Taxes and Investment Tax Credit (txdtc)</p> <p>Liabilities - Other - Total (lo)</p> <p>Liabilities - Total (lt)</p> <p>Noncontrolling Interest (Balance Sheet) (mib)</p> <p>Preferred/Preference Stock (Capital) - Total (pstk)</p> <p>Common/Ordinary Equity - Total (ceq)</p> <p>Stockholders Equity - Parent (seq)</p> <p><u>Statement of Cash Flows Variables - Investing Activities</u></p> <p>Increase in Investments (ivch)</p> <p>Sale of Investments (siv)</p> <p>Short-Term Investments - Change (ivstch)</p> <p>Capital Expenditures (capx)</p> <p>Sale of Property (sppe)</p> <p>Acquisitions (aqc)</p> <p>Investing Activities - Other (ivaco)</p> <p>Investing Activities - Net Cash Flow (ivncf)</p>	<p><u>Income Statement Variables</u></p> <p>Sales/Turnover (Net) (sale)</p> <p>Cost of Goods Sold (cogs)</p> <p>Selling, General and Administrative Expense (xsga)</p> <p>Operating Income Before Depreciation (oibdp)</p> <p>Operating Income After Depreciation (oiadp)</p> <p>Interest and Related Expense - Total (xint)</p> <p>Nonoperating Income (Expense) (nopi)</p> <p>Special Items (spi)</p> <p>Pretax Income (pi)</p> <p>Income Taxes - Total (txt)</p> <p>Income Before Extraordinary Items (ib)</p> <p>Net Income Adjusted for Common/Ordinary Stock (Capital)</p> <p>Equivalents (niadj)</p> <p>Earnings Per Share (Basic) - Excluding Extraordinary Items (epspx)</p> <p>Earnings Per Share (Diluted) - Excluding Extraordinary Items (epsfx)</p> <p><u>Statement of Cash Flows Variables - Operating Activities</u></p> <p>Income Before Extraordinary Items (Cash Flow) (ibc)</p> <p>Extraordinary Items and Discontinued Operations (Cash Flow) (xidoc)</p> <p>Depreciation and Amortization (Cash Flow) (dpc)</p> <p>Deferred Taxes (Cash Flow) (txdc)</p> <p>Equity in Earnings - Unconsolidated Subsidiaries (esub)</p> <p>Sale of Property, Plant and Equipment and Investments - Gain (Loss) (sppiv)</p> <p>Funds from Operations - Other (fopo)</p> <p>Accounts Receivable - Decrease (Increase) (recch)</p> <p>Inventory - Decrease (Increase) (invch)</p> <p>Accounts Payable and Accrued Liabilities - Increase/(Decrease) (apalch)</p> <p>Income Taxes - Accrued - Increase/(Decrease) (txach)</p> <p>Assets and Liabilities - Other - Net Change (aoloch)</p> <p>Operating Activities - Net Cash Flow (oancf)</p> <p><u>Statement of Cash Flows Variables - Financing Activities</u></p> <p>Sale of Common and Preferred Stock (sstk)Excess Tax Benefit of Stock Options - Cash Flow Financing (txbcof)Purchase of Common and Preferred Stock (prstk)Cash Dividends (Cash Flow) (dv)Long-Term Debt - Issuance (dltis)Long-Term Debt - Reduction (dltr)Current Debt - Changes (dlcch)Financing Activities - Other (fiao)Financing Activities - Net Cash Flow (fincf)</p>
--	--

Figure 28: Description of each column in the Annual Simplified Financial Statement.