



TableVista: Benchmarking Multimodal Table Reasoning under Visual and Structural Complexity

Zheyuan Yang^{1*} Liqiang Shang^{2*} Junjie Chen³ Xun Yang¹ Chenglong Xu¹
Bo Yuan¹ Chenyuan Jiao¹ Yaoru Sun^{1†} Yilun Zhao⁴

¹Tongji University ²University of Bristol ³Tianjin University ⁴Yale University

Abstract

We introduce TableVista, a comprehensive benchmark for evaluating foundation models in multimodal table reasoning under visual and structural complexity. TableVista consists of 3,000 high-quality table reasoning problems, where each instance is expanded into 10 distinct visual variants through our multi-style rendering and transformation pipeline. This process encompasses diverse scenario styles, robustness perturbations, and vision-only configurations, culminating in 30,000 multimodal samples for a multi-dimensional evaluation. We conduct an extensive evaluation of 29 state-of-the-art open-source and proprietary foundation models on TableVista. Through comprehensive quantitative and qualitative analysis, we find that while evaluated models remain largely stable across diverse rendering styles, they exhibit pronounced performance degradation on complex structural layouts and vision-only settings, revealing that current models struggle to maintain reasoning consistency when structural complexity combines with visually integrated presentations. These findings highlight critical gaps in current multimodal capabilities, providing insights for advancing more robust and reliable table understanding models.

 **Data** [TableVista](#)
 **Code** [FlowRays/TableVista](#)

1 Introduction

Tables serve as the bedrock of structured data across diverse professional domains, yet the optimal way for foundation models to perceive and reason over them remains an open question (Deng et al., 2024; Zhou et al., 2025b; Zhao et al., 2024b). Traditionally, table reasoning has relied on text-based serialization (*e.g.*, Markdown or HTML),

but this approach often collapses the rich, spatial-structural information inherent in complex layouts (Zhang et al., 2023; Yang et al., 2025b). In contrast, multimodal table reasoning (*i.e.*, treating tables as visual artifacts) naturally preserves intricate features such as multi-level headers, merged cells, and multi-table alignments that are frequently mangled in raw text strings (Zheng et al., 2024; Zhao et al., 2024a; Zhou et al., 2025a). This visual framing also better reflects real user interactions, in which tables are typically encountered as rendered documents or interface elements rather than as linearized data. However, the extent to which current multimodal foundation models can sustain consistent reasoning across diverse visual presentations remains underexplored, especially under structural complexity and real-world perturbations.

To bridge this gap and provide a rigorous assessment of multimodal table reasoning, we introduce TableVista, a comprehensive benchmark designed to evaluate foundation models under conditions of high visual and structural complexity. Unlike existing datasets that often rely on simplistic or uniform table renderings, TableVista comprises 3,000 high-quality reasoning problems, each expanded into 10 distinct visual variants through a sophisticated multi-style rendering and transformation pipeline. This pipeline systematically introduces diverse scenario styles, robustness perturbations, and real-world environmental perturbations, resulting in a large-scale collection of 30,000 multimodal samples. By providing such a multi-dimensional evaluation suite, TableVista forces models to go beyond basic OCR-like recognition and demands a deeper synthesis of visual perception and logical reasoning across heterogeneous presentations.

We conduct an extensive evaluation of 29 state-of-the-art foundation models, including both proprietary and open-source models. We find that performance degrades sharply as structural complexity and reasoning depth increase, with open-

*Equal contributions. †Correspondence to: Yaoru Sun (yaoru@tongji.edu.cn)

Benchmark	Multimodal	Structural Complexity			Visual Diversity & Robustness		
		Hierarch.	Long	Multi	Scenario	Transformation	Vision-Only
HiTab (Cheng et al., 2021)	✗	✓	✗	✗	–	–	–
NeedleInATable (Wang et al., 2025a)	✗	✓	✓	✗	–	–	–
MMQA (Wu et al., 2025)	✗	✗	✓	✓	–	–	–
TableVQA-Bench (Kim et al., 2024)	✓	✗	✗	✗	✗	✗	✗
MMTabQA (Mathur et al., 2024)	✓	✗	✗	✗	✗	✗	✗
MMTab (Zheng et al., 2024)	✓	✓	✓	✗	✗	✗	✗
MMTBench (Titiya et al., 2025)	✓	✓	✓	✗	✗	✗	✗
TableVista (ours)	✓	✓	✓	✓	✓	✓	✓

Table 1: Comparison of **TableVista** with representative table understanding benchmarks. TableVista is the first benchmark to integrate diverse structural complexities with a rigorous multi-style visual robustness pipeline.

source models exhibiting more severe collapse. Beyond reasoning, visual conditions introduce an additional failure mode: while models remain stable across rendering themes, they are brittle under realistic perturbations such as fragmented layouts and phone-captured tables, where errors are driven more by spatial misalignment than text recognition. Prompting further reveals a stark contrast—chain-of-thought largely narrows the gap between open-source and proprietary models, yet substantial disparities re-emerge under direct-output prompting, indicating limited internalization of multi-step tabular reasoning. These findings highlight critical bottlenecks in current multimodal models, particularly in their ability to robustly anchor structured reasoning in visual space.

2 Related Work

Table Reasoning Evaluation. Early evaluation benchmarks for table reasoning primarily focused on text-only serialization, addressing tasks like question answering (Pasupat and Liang, 2015; Nan et al., 2021; Zhao et al., 2023a,b) and fact verification (Chen et al., 2019). However, these unimodal benchmarks overlook the rich spatial structure inherent in visual layouts. Recent benchmarks have begun to incorporate visual modalities. For instance, TableVQA-Bench (Kim et al., 2024) and MMTab (Zheng et al., 2024) utilize synthetic rendering pipelines to evaluate visual reasoning capabilities. To account for richer table content and visual elements, MMTabQA (Mathur et al., 2024) introduces knowledge-aware reasoning over tables containing symbols and icons, while MMTBench (Titiya et al., 2025) incorporates interleaved charts and images. In specific domains, MMSci (Yang et al., 2025a) targets numerical reasoning within scientific tables. Most recently,

TABLET (Alonso et al., 2025) has emphasized large-scale robust evaluation on tables rendered from original web sources. However, as shown in Table 1, existing benchmarks typically present tables in a fixed, ideal, or single original visual state. They neglect the critical dimension of reasoning consistency under visual perturbations, where a model should yield identical answers despite changes in layout, style, or orientation.

Multimodal Table Reasoning. Recent progress in multimodal table understanding has been driven by MLLMs, which enable direct reasoning over rendered tables and other visually grounded document elements. Pioneering generalist models such as Table-LLaVA (Zheng et al., 2024) and TabPedia (Zhao et al., 2024a) enhance performance through table-specific instruction tuning and concept synergy mechanisms. To address specific limitations in data and resolution, SynTab-LLaVA (Zhou et al., 2025a) introduces a decoupled synthesis pipeline to scale up training data efficiently, while MMSci (Yang et al., 2025a) employs dynamic input resolution to capture fine-grained details in scientific tables. More recently, research has pivoted toward advanced architectures and reasoning strategies. TableMoE (Zhang et al., 2025) proposes a neuro-symbolic routing mechanism within a Mixture-of-Experts (MoE) framework to handle structurally degraded “WildStruct” tables. To bridge the modality gap, Turbo (Jiang et al., 2025) leverages privileged structured information during training to distill reasoning capabilities from strong reasoning models such as DeepSeek-R1 (DeepSeek-AI et al., 2025).

3 TableVista Evaluation Benchmark

This section details the construction process of the TableVista benchmark. The pipeline is designed to

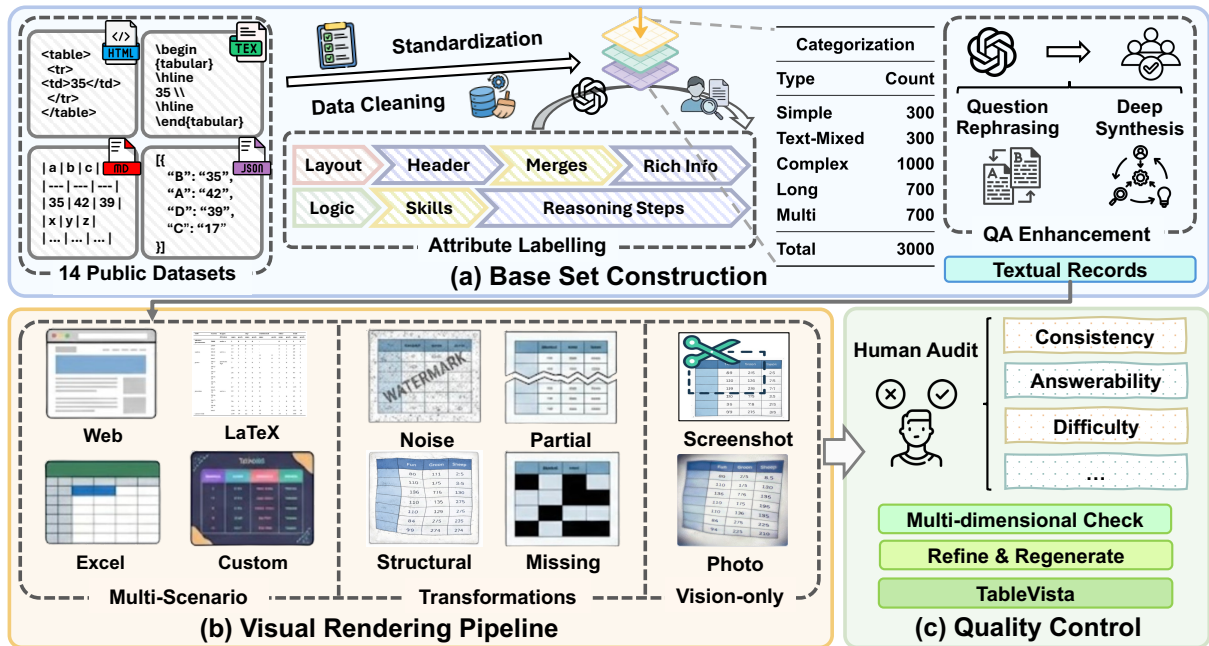


Figure 1: An overview of the TableVista benchmark construction pipeline.

establish a rigorous benchmark for evaluating the table reasoning capabilities of MLLMs by integrating diverse table structural complexities with challenging visual transformations. Figure 1 provides an overview of benchmark construction workflow.

3.1 Base Set Construction

We first construct a textual **Base Set** by collecting and re-annotating samples from public table reasoning datasets. We engage 12 expert annotators (detailed information provided in Appendix A) to re-annotate these samples with the assistance of state-of-the-art LLMs¹. The Base Set construction process is as follows:

Data Standardization and Attribute Labeling.

We first aggregate raw records from 14 prominent public datasets (e.g., WTQ (Pasupat and Liang, 2015), HiTab (Cheng et al., 2021), TabFact (Chen et al., 2019), MMQA (Wu et al., 2025), FinQA (Chen et al., 2021), etc.), standardizing them into a canonical format comprising the *table*, *textual context*, *question*, and *answer*. For each record, we perform structured attribute labeling consisting of four tasks: (1) identifying table layout attributes (e.g., multi-level headers, merged cells, composite structures) and recording dimensional information, (2) assigning an *information richness* score (1–5) based on data diversity, relation den-

sity, and numerical computability, (3) characterizing the reasoning requirements of each QA pair by assigning *skill scores* across four dimensions—lookup, aggregation, numerical, and logical—and (4) counting the total *reasoning steps* required to derive the answer. These labels are initially produced by GPT-5, and the annotators then verify each result by inspecting the model’s chain-of-thought, re-annotating any cases they find incorrect.

Taxonomic Categorization and Multidimensional Filtering.

Based on the labeled attributes, we perform rigorous filtering across two dimensions: table structural complexity and QA reasoning difficulty. First, we categorize records into five structural types with explicit quotas: *Simple Structure* (standard grid), *Text-Mixed* (with associated textual context), *Complex Structure* (with hierarchical headers or merged cells), *Long Tables* (with large row or column dimensions), and *Multi-Table* scenarios (involving 2–5 associated tables). Within each structural quota, we prioritize samples with the highest reasoning potential. The selection criteria emphasize higher table *information richness* scores, higher aggregate *skill scores*, and larger numbers of *reasoning steps*. Through this hierarchical selection procedure, we ultimately retain 3,000 highly representative and reasoning-intensive records, yielding a benchmark that balances structural diversity and reasoning capability.

¹Unless otherwise specified, we use GPT-5 by default throughout the data construction process.

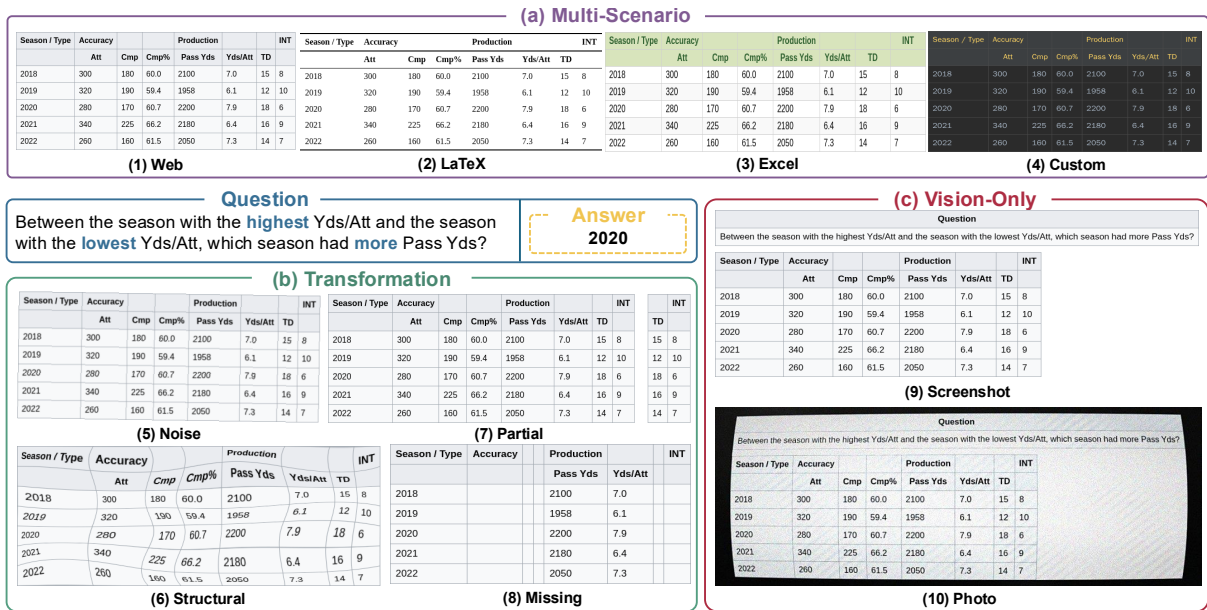


Figure 2: Samples of TableVista. The central box illustrates a standard question–answer pair requiring calculation and reasoning over the table. (a) Multi-Scenario: four rendering styles—Web, LaTeX, Excel, and Custom. (b) Transformation: four robustness perturbations—Noise, Structural, Partial, and Missing. (c) Vision-Only: two settings simulating real-world capture—Screenshot and Photo.

QA Enhancement.

For the filtered samples, annotators further enhance QA pairs using a dual-track strategy: (1) *Question Rephrasing*: if the original QA already meets the target reasoning complexity, annotators rewrite the question to increase linguistic diversity and reduce potential contamination while preserving the answer; and (2) *Deep Synthesis*: if the original QA is trivial, annotators author a new multi-step reasoning question grounded in the table and textual context, requiring at least two core skills and multiple logical hops. To ensure evaluative precision, all answers are standardized as unique short phrases or numbers.

3.2 Multi-Style Visual Pipeline

The primary contribution of TableVista is its sophisticated visual rendering pipeline, which systematically maps textual foundations onto a broad spectrum of visual representations, spanning diverse stylistic scenarios, structural and environmental perturbations, and unified vision-only configurations. This multi-faceted framework enables a rigorous evaluation of model robustness by simulating the complex interplay between structural layout, aesthetic variation, and simulated visual degradation, as illustrated in Figure 2.

Multi-Scenario Rendering. To capture the diversity of table formats encountered in practice, we simulate four authentic scenarios that reflect dis-

tinct data ecosystems. The *Web* scenario replicates the standard HTML layouts common in Wikipedia, while *LaTeX* mimics the professional typography found in academic literature. The *Excel* scenario reproduces the dense, grid-centric interfaces characteristic of spreadsheet applications. Finally, the *Customized* scenario introduces high-variance visual noise by sampling from ten predefined thematic styles featuring diverse font families and color palettes. All scenarios are instantiated via a headless browser where viewports are dynamically optimized using localized bounding box metadata. This ensures that the resulting high-resolution captures maintain structural integrity without unintended cropping.

Robustness Transformations.

Beyond stylistic variation, we introduce a suite of transformations to simulate the degraded or complex conditions often found in “in-the-wild” capture environments. Using the *Web* scenario as a baseline, we apply four distinct perturbations. *Noise* transformations incorporate stochastic combinations of Gaussian speckle, resolution downsampling, rotational tilt, and translucent watermarks. *Structural Noise* transformations introduce structure-specific degradations, such as faded gridlines and geometric warping, that disrupt the table’s visual skeleton while preserving its textual content. The *Partial* trans-

Metric	Count / Value
<i>Dataset Size</i>	
Tables	4,449
QA Pairs	3,000
Visual Samples	30,000
<i>Structural Distribution</i>	
Simple Structure	300
Text-Mixed	300
Complex Structure	1,000
Long Table	700
Multi-Table	700
<i>Content Statistics</i>	
Avg. Question Words	26.2
Avg. Answer Words	1.4
Avg. Table Rows	15.3
Avg. Table Columns	6.6
<i>Reasoning Complexity (Steps / Skill Score)</i>	
Easy	3.7 / 11.2
Medium	5.2 / 13.8
Hard	6.9 / 15.1

Table 2: TableVista data statistics.

formation segments the table into disjoint visual blocks along structural boundaries to test long-range spatial reasoning. Lastly, *Missing* transformations mask arbitrary cell contents without row- or column-level restrictions, while annotators verify that the remaining cells retain sufficient information to answer the target query.

Vision-Only Setting and Realism Artifacts. To align the evaluation with naturalistic user interactions, we propose a *Vision-Only* configuration where the model processes a single image without separate text input. This setting is operationalized through two distinct modalities. The *Screen Capture* mode renders the question, context, and table into a single unified interface, while the *Simulated Photo* mode applies synthetic camera artifacts, such as moiré patterns and perspective distortion, directly on top of the *Screen Capture* images to emulate the visual degradation of mobile photography. These challenges necessitate a synergistic capability to extract questions and interpret tabular data within a purely visual domain.

3.3 Quality Control and Human Audit

We conduct a *human-in-the-loop* audit throughout the pipeline, covering both the 3,000 textual base samples and the 30,000 rendered images produced by the visual pipeline. Expert reviewers perform a multi-dimensional check across consistency, answerability, difficulty, and other quality attributes,

with particular attention to ensuring answerability even after severe visual transformations such as masking or splitting. Samples that fail any dimension are refined or regenerated until they meet the required standard. The resulting benchmark, TableVista, demonstrates high logical rigor and scenario diversity, providing a superior quality benchmark for multimodal table understanding. Detailed rejection and revision statistics at each stage of the pipeline are reported in Appendix B.

4 Experiments

In this section, we first describe the experimental setup and then provide a detailed analysis of the experimental results.

4.1 Experimental Setup

Evaluated Models. We evaluate the performance of state-of-the-art foundation models on our benchmark, including both proprietary and open-source models spanning a wide range of scales and reasoning capabilities. Proprietary models evaluated are GPT-5.4 (OpenAI, 2026a) and GPT-5.4-mini (OpenAI, 2026b). Open-source models include Qwen2.5-VL (Bai et al., 2025b), Qwen3-VL (Bai et al., 2025a), Qwen3.5 (Team, 2026), Qwen3.6 (Qwen Team, 2026), InternVL3.5 (Wang et al., 2025b), Llama-4 (Meta, 2025), Gemma-3 (Kamath et al., 2025), Gemma-4 (DeepMind, 2026), Molmo2-8B (for AI, 2025), MiniCPM-V-4.5 (Yu et al., 2025), LLaVA-v1.5 (Liu et al., 2023), LLaVA-v1.6 (Liu et al., 2024), and TableLLaVA (Zheng et al., 2024). Detailed model specifications and inference configurations are provided in Appendix C.1.

Evaluation Metrics. We employ a hybrid evaluation strategy combining Exact Match (EM) and an LLM-based Judge. We primarily assess correctness via normalized string matching. To mitigate penalties from minor formatting deviations, we utilize GPT-5-mini to verify semantic consistency for samples failing the EM criterion.

Evaluation Prompts. For the analyses in §4.2 and §4.3, we evaluate all models under a *direct-output* prompt without enabling thinking mode (for GPT-5.4 and GPT-5.4-mini, we correspondingly set `reasoning_effort` to `none`). We adopt this protocol because our preliminary experiments reveal that explicit reasoning compresses the performance landscape: under a CoT prompt

Model	Structural Type					Reasoning Difficulty			Overall
	Simple	Text-Mixed	Complex	Long	Multi	Easy	Medium	Hard	
<i>Proprietary Models</i>									
GPT-5.4	73.0	86.7	81.7	68.9	61.3	93.6	80.1	47.0	73.6
GPT-5.4-mini	52.0	<u>61.0</u>	59.3	48.7	40.0	64.0	<u>56.1</u>	<u>35.2</u>	51.8
<i>Open-Source Models</i>									
Qwen2.5-VL-72B	52.3	58.3	<u>59.7</u>	52.1	<u>53.1</u>	<u>90.2</u>	54.3	22.1	<u>55.5</u>
Gemma-4-31B-it	<u>57.3</u>	54.0	<u>57.6</u>	<u>54.4</u>	<u>52.3</u>	<u>88.2</u>	55.6	21.9	<u>55.2</u>
Llama-4-Maverick	55.3	55.7	55.9	52.3	52.4	84.4	53.7	24.4	54.2
Qwen2.5-VL-32B	51.0	49.7	54.5	48.3	49.1	84.5	49.0	19.4	51.0
Qwen3.5-27B	50.3	45.3	53.0	50.9	48.6	84.2	47.5	19.6	50.4
Llama-4-Scout	50.7	47.7	49.6	49.7	48.1	81.8	47.0	18.8	49.2
Qwen3.5-122B-A10B	45.0	48.3	52.2	49.4	46.3	82.6	45.7	18.9	49.1
InternVL3.5-14B	44.3	48.0	48.9	46.4	41.7	83.0	42.1	13.2	46.1
Gemma-4-26B-A4B-it	45.3	45.7	50.1	44.1	42.0	79.7	41.4	16.6	45.9
InternVL3.5-30B-A3B	46.7	49.0	47.7	43.3	43.3	81.6	41.3	14.1	45.7
Qwen3-VL-30B-A3B	44.3	48.7	46.8	42.4	41.0	78.9	39.4	14.8	44.4
Qwen2.5-VL-7B	43.0	41.3	47.5	41.9	42.4	79.0	38.6	14.2	43.9
Gemma-3-27B-it	46.7	44.0	44.6	43.7	39.6	76.5	39.5	14.1	43.4
Qwen3-VL-8B	40.7	44.0	44.1	41.9	39.9	76.7	37.3	12.7	42.2
Molmo2-8B	39.0	40.7	45.2	38.9	39.4	75.9	36.0	12.0	41.3
Qwen3.6-35B-A3B	41.3	38.0	42.5	40.3	41.0	70.9	35.9	16.4	41.1
Qwen3.5-35B-A3B	44.0	38.0	42.9	38.6	38.3	70.6	37.1	13.6	40.4
Qwen3.5-9B	42.3	41.3	40.3	39.1	39.4	72.6	33.7	14.1	40.1
InternVL3.5-8B	35.7	43.0	40.4	38.4	39.1	73.9	34.5	9.9	39.4
Gemma-3-12B-it	39.7	37.3	42.0	36.7	34.4	70.6	32.8	11.5	38.3
MiniCPM-V-4.5	33.3	34.3	37.4	33.3	33.1	63.5	28.3	12.4	34.7
Gemma-4-E4B-it	30.3	23.7	31.0	26.3	26.1	53.1	23.2	7.6	28.0
LLaVA-v1.6-Vicuna-13B	16.0	17.0	18.6	15.1	14.9	30.8	12.8	5.9	16.5
LLaVA-v1.6-Vicuna-7B	14.0	16.3	16.9	13.4	14.6	27.1	13.2	5.3	15.2
Table-LLaVA-v1.5-13B	13.3	11.0	11.9	12.7	11.7	20.6	10.4	5.3	12.1
Table-LLaVA-v1.5-7B	11.0	11.0	7.8	9.4	9.3	16.7	6.8	4.0	9.2
LLaVA-v1.5-7B	6.0	5.7	6.6	6.4	6.7	10.6	4.7	4.0	6.4

Table 3: Model performance in the *Web* scenario, stratified by structural complexity and reasoning difficulty.

with thinking mode enabled, even Qwen3-VL-8B-Thinking reaches 86.0% on TableVista, narrowing its gap with both proprietary SOTA and substantially larger open-source SOTA models to a few percentage points. In contrast, under the direct-output setting Qwen3-VL-8B-Instruct attains only 43.8%, and pronounced gaps re-emerge across proprietary SOTA, open-source SOTA, and smaller open-source models. We therefore first analyze model performance along structural, reasoning, and visual-robustness dimensions under the direct-output setting, and subsequently provide a dedicated analysis of CoT-versus-direct-output behavior in §4.4. The full prompt templates for both settings are provided in Appendix C.2.

4.2 Evaluation on Table Structural and Reasoning Complexity

Table 3 reports the performance breakdown across five structural categories and three reasoning difficulty levels. Our primary findings are as follows:

Structural Complexity Challenges. Complex structural layouts pose substantial challenges for current MLLMs, though the pattern diverges from simple intuition. For the frontier model GPT-5.4, accuracy follows the expected structural hierarchy, peaking on Text-Mixed (86.7%) and Complex (81.7%) tables and dropping sharply on Long (68.9%) and Multi-table (61.3%) scenarios. For weaker models, however, even Simple grids remain challenging (*e.g.*, Qwen2.5-VL-72B at 52.3%), since our QA pairs demand intensive reasoning that a simple layout alone cannot offset. The Multi-table scenario proves uniquely difficult across the board, with most open-source models failing to exceed 50%, suggesting that TableVista requires spatial reasoning capabilities that remain difficult for current models, especially when context must be maintained across disjoint regions.

Reasoning Intensity and Table Skills. Performance varies considerably across reasoning difficulties mirroring the skill stratification defined in

Model	Scenario Style				Visual Transformation				Vision-Only		Avg.
	Web	LaTeX	Excel	Custom	Noise	Structural	Partial	Missing	Screenshot	Photo	
<i>Proprietary Models</i>											
GPT-5.4	73.6	72.2	71.9	72.0	70.8	70.4	68.8	84.8	69.4	67.3	72.1
GPT-5.4-mini	51.8	49.9	50.1	51.1	49.5	48.4	46.8	66.9	42.0	37.7	49.4
<i>Open-Source Models</i>											
Qwen2.5-VL-72B	55.5	54.5	54.5	55.0	51.1	54.4	50.5	71.3	57.4	54.0	55.8
Llama-4-Maverick	54.2	53.2	53.9	52.7	53.5	54.4	52.9	66.9	53.4	51.8	54.7
Gemma-4-31B-it	55.2	56.5	55.9	54.4	54.3	53.3	51.2	75.7	42.2	44.2	54.3
Qwen2.5-VL-32B	51.0	50.7	50.7	51.9	49.3	49.1	48.1	67.3	51.1	49.8	51.9
Qwen3.5-27B	50.4	49.4	48.3	50.2	49.2	47.3	46.7	63.8	55.1	53.5	51.4
Qwen3.5-122B-A10B	49.1	48.5	48.6	48.9	49.0	48.7	47.3	63.3	54.3	53.5	51.1
Llama-4-Scout	49.2	48.6	47.7	48.1	47.5	48.3	46.0	64.9	46.7	47.2	49.4
InternVL3.5-14B	46.1	45.5	45.1	45.4	45.3	45.2	43.2	63.0	42.7	41.6	46.3
InternVL3.5-30B-A3B	45.7	44.4	44.4	45.6	45.6	44.9	43.5	65.1	40.1	35.5	45.5
Qwen2.5-VL-7B	43.9	44.3	44.4	44.1	42.5	42.1	41.0	61.8	41.6	40.0	44.6
Qwen3-VL-30B-A3B	44.4	42.7	42.7	45.4	43.0	41.8	40.4	60.2	41.1	42.4	44.4
Qwen3-VL-8B	42.2	41.8	41.7	42.0	41.4	41.9	39.2	57.7	45.3	44.6	43.8
Qwen3.6-35B-A3B	41.1	41.5	41.6	42.1	41.6	39.9	38.9	56.6	46.1	44.5	43.4
Gemma-3-27B-it	43.4	42.7	42.2	40.9	40.1	41.0	39.0	61.1	40.1	37.5	42.8
Gemma-4-26B-A4B-it	45.9	46.8	44.4	46.0	43.2	43.4	39.9	65.6	29.4	20.7	42.5
Qwen3.5-9B	40.1	39.2	39.1	39.4	39.5	39.3	37.6	55.2	39.7	37.9	40.7
InternVL3.5-8B	39.4	40.6	39.4	39.8	39.2	39.0	37.1	58.5	37.1	36.5	40.7
Qwen3.5-35B-A3B	40.4	38.9	38.1	38.3	38.4	37.9	36.8	54.9	39.7	37.0	40.0
Molmo2-8B	41.3	40.8	42.2	41.7	37.0	38.8	37.0	57.5	33.2	29.7	39.9
Gemma-3-12B-it	38.3	38.5	38.6	38.7	36.2	37.2	35.7	55.0	33.9	32.9	38.5
MiniCPM-V-4.5	34.7	31.8	33.3	34.9	34.0	32.8	32.4	52.7	28.0	27.5	34.2
Gemma-4-E4B-it	28.0	29.9	27.0	26.4	26.7	27.8	24.8	47.6	16.9	15.3	27.0
LLaVA-v1.6-Vicuna-13B	16.5	16.1	16.2	16.2	16.3	16.2	15.0	31.4	2.9	2.8	14.9
LLaVA-v1.6-Vicuna-7B	15.2	15.3	15.0	14.5	14.7	14.9	14.1	28.2	3.8	3.0	13.9
Table-LLaVA-v1.5-13B	12.1	12.2	11.0	12.5	12.1	12.6	11.6	16.8	1.6	1.8	10.4
Table-LLaVA-v1.5-7B	9.2	9.1	7.5	10.1	9.2	9.4	8.8	13.8	0.9	1.1	7.9
LLaVA-v1.5-7B	6.4	6.2	6.2	7.0	6.9	6.7	6.4	10.4	0.5	0.4	5.7

Table 4: Model performance under different real-world scenarios, visual transformations, and vision-only settings.

§3. GPT-5.4 achieves 93.6% on Easy tasks but attains just 47.0% on Hard tasks requiring multi-hop aggregation. The gap widens further for open-source models, with even the strongest Qwen2.5-VL-72B collapsing from 90.2% to 22.1%. This uneven performance confirms that solving visual grounding does not automatically equate to solving complex tabular logic and highlights the need for future models to improve reasoning depth.

Proprietary vs. Open-source Models. Open-source foundation models remain notably behind proprietary models, with a substantial gap that persists across evaluations. Across all evaluated tasks, proprietary models consistently outperform open-source models with GPT-5.4 achieving 73.6% overall accuracy. The top open-source model Qwen2.5-VL-72B achieves 55.5% and trails GPT-5.4 by roughly 18 percentage points. However, this model significantly outperforms older architectures like LLaVA-v1.5-7B which fails to reach 7% accuracy. This highlights continued progress in open-source development yet also points to substantial remain-

ing challenges in achieving robust multimodal table reasoning.

4.3 Evaluation on Real-World Scenario and Visual Transformation

To assess the robustness of MLLMs against distribution shifts, we evaluate performance under diverse rendering styles and visual perturbations. The experiment results are presented in Table 4.

Scenario-Specific Analysis. Models exhibit remarkable consistency across different table scenarios. Contrary to the structural sensitivity observed in previous sections, performance remains surprisingly stable across *Web*, *LaTeX*, *Excel*, and *Customized* themes. For instance, GPT-5.4 maintains an accuracy around 72% across all four styles, and the leading open-source model Qwen2.5-VL-72B shows similar stability with scores hovering around 55%. This uniformity suggests that current foundation models have successfully generalized to diverse visual aesthetics and font variations.

Visual Transformation Sensitivity. Visual perturbations reveal distinct vulnerability patterns in model robustness. Counter-intuitively, the *Missing* setting, which masks arbitrary cells while preserving answer-sufficient content, yields the highest scores across nearly all models—GPT-5.4 rises to 84.8% and Qwen2.5-VL-72B reaches 71.3%. This indicates that reducing visible clutter can actually concentrate model attention onto reasoning-critical cells. In contrast, the *Partial* setting, which fragments the table into disjoint visual blocks, proves the most challenging overall, with GPT-5.4 dropping to 68.8% and most open-source models falling below 50%. *Noise* and *Structural Noise* perturbations cause moderate and comparable degradation, suggesting that current models tolerate low-level visual noise reasonably well but struggle to maintain spatial coherence once the table’s structural continuity is broken.

Vision-Only Setting Analysis. The *Vision-Only* setting reveals how models cope when questions and table content must be jointly parsed from a single image. Under the *Screen Capture* mode, capable models largely retain their baseline competence—GPT-5.4 scores 69.4% and Qwen2.5-VL-72B reaches 57.4%, comparable to their performance with separately provided inputs. The *Simulated Photo* mode, which overlays synthetic moiré patterns and perspective distortion on the rendered interface, induces further degradation: GPT-5.4 drops modestly to 67.3%, while mid-sized models deteriorate more sharply (e.g., Gemma-4-26B-A4B-it falls from 29.4% to 20.7%). A distinct failure mode emerges for legacy architectures such as the LLaVA family, whose accuracy collapses to near-zero under both vision-only modes, indicating an inability to extract questions and tabular content simultaneously from a unified visual input. Together, these results expose two separate bottlenecks: robustness to camera-style artifacts for capable models, and the more fundamental challenge of synergistic visual perception for weaker ones.

4.4 CoT vs. Direct-Output Prompting

Building on the DO-based analyses in §4.2 and §4.3, we now quantify the landscape-compression effect noted in §4.1. Table 5 contrasts direct-output (DO) and chain-of-thought (CoT) prompting for five representative models. All models gain from explicit reasoning, but the gains are inversely tied

Model	DO	CoT	Δ
GPT-5.4	72.1	95.6	+23.5
GPT-5.4-mini	49.4	91.5	+42.1
Qwen3.5-27B	51.4	96.2	+44.8
Gemma-4-31B-it	54.3	86.1	+31.8
Qwen3-VL-8B [†]	43.8	86.0	+42.2

Table 5: Overall accuracy (%) under direct-output (DO) and CoT prompting. [†]: Direct output uses Qwen3-VL-8B-Instruct; CoT uses Qwen3-VL-8B-Thinking.

to DO strength: GPT-5.4 rises only +23.5 points (72.1 to 95.6), while weaker DO performers close most of the gap (GPT-5.4-mini +42.1, Qwen3.5-27B +44.8, Gemma-4-31B-it +31.8). The effect is most striking at small scale: Qwen3-VL-8B jumps from 43.8 to 86.0 (+42.2), nearly erasing the gap to far larger models under CoT. The DO bottleneck on TableVista thus lies less in visual perception, which §4.3 shows to be handled well across rendering styles, than in packing multi-step tabular computation into a single forward pass. CoT also reshuffles the leaderboard: Qwen3.5-27B overtakes GPT-5.4 at the top, and the Gemma-Qwen ordering flips by over ten points. These contrasts position DO and CoT as complementary rather than redundant: CoT approaches the ceiling and reveals whether the required capabilities are present when reasoning is externalized, while DO stresses whether they can be internalized end-to-end and still leaves substantial headroom across all evaluated models.

4.5 Qualitative Analysis

We conduct a comprehensive qualitative audit to gain deeper insight into failure modes that aggregate metrics do not capture. Specifically, we randomly sample 100 examples from the error cases produced by GPT-5.4 and Qwen3.5-27B. We then manually analyze these examples to characterize the underlying failures and identify three primary failure modes:

Visual Brittleness. This case highlights incomplete visual recognition under noisy input as the primary failure mode. As shown in Figure 7, GPT-5.4 under clean rendering correctly executes a Filter \rightarrow Discriminate \rightarrow Select pipeline and identifies *Product-differentiated* ($\Delta = -17$ pp vs. Scale-based’s -12 pp). Under simulated photo, however, degradation in the upper image region leads to partial question understanding: while the model correctly extracts all relevant numerical values, it

fails to recognize the superlative qualifier *largest*. This selective omission causes the model to operate on incomplete information—after identifying sectors satisfying the base condition, its internal objective is effectively reduced from *Filter-and-Argmax* to *Filter-and-Verify*, resulting in the output {Scale-based, Product-differentiated}. The error thus reflects localized visual brittleness: missing fine-grained textual cues leads to answers that are internally consistent with perceived inputs, yet incomplete with respect to the full query, without any calibration signal to indicate the mismatch.

Spatial Alignment Failures. Structural complexity places substantial demands on spatial reasoning, especially in dense tables with multi-level headers. As shown in Figure 8, a representative *cross-row argmin* error arises at the filtering stage. Under clean rendering, GPT-5.4 correctly compares *Men* = 15.2 and *Women* = 11.8, selects *Women*, and returns 22.9. Under simulated photo conditions, however, the reasoning trace remains logically consistent while perception fails: 11.8 is misread as 15.8 due to a 1 → 5 substitution influenced by the nearby 15.2. This misperception flips the comparison, leading the model to select *Men* and retrieve 16.3 pp, a valid but incorrect cell. Notably, row alignment remains intact; the failure originates from fine-grained digit confusion, which propagates to the decision stage while remaining undetected in the CoT. This highlights a key limitation: coarse structural grounding is insufficient, as models struggle to distinguish visually similar digits within correctly localized cells without stronger sub-cell disambiguation.

Error Distribution. To pinpoint failure causes, we analyzed the error distribution, as shown in Figure 3. *Table Understanding* dominates the error distribution (54%), significantly surpassing *Reasoning & Calculation* (29%) and *Visual Perception* (12%). Within table understanding, *Spatial Alignment* (32%) and *Structure Parsing* (22%) are the primary bottlenecks. Notably, errors stemming solely from text recognition are rare (4%), confirming that the challenge lies in maintaining spatial structure rather than basic perception.

5 Conclusion

We present TableVista, a comprehensive benchmark designed to rigorously evaluate the multimodal table reasoning capabilities of foundation

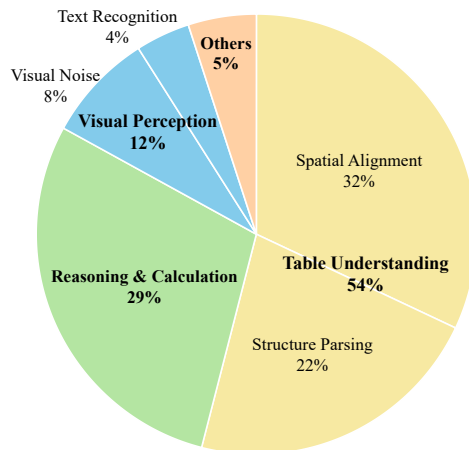


Figure 3: Error distribution analysis on TableVista. The majority of failures stem from Table Understanding (54%), followed by Reasoning & Calculation (29%) and Visual Perception (12%).

models under visual and structural complexity. Each instance in TableVista is expanded into distinct visual variants through a rigorous multi-style rendering pipeline and audited by human experts. Our evaluation of 29 frontier foundation models reveals that even the most advanced proprietary models degrade sharply as structural and reasoning complexity grow, remain brittle under realistic visual conditions, and show pronounced open-source to proprietary gaps under direct-output prompting. These findings expose critical limits in anchoring structured reasoning in visual space, providing valuable insights for advancing the robust and structure-aware reasoning abilities of future foundation models.

Limitations

This work focuses on evaluating frontier models for robust multimodal table understanding. We do not explore or propose new training methods to directly mitigate the structural alignment failures identified in our tasks. Future research could build upon our analyses to investigate effective visual instruction tuning strategies that enhance spatial grounding in complex document layouts. Moreover, beyond the tabular data covered in our benchmark, real-world documents often involve interleaved modalities such as charts and natural images. We therefore encourage future work to extend our study by incorporating a broader range of visual document elements and heterogeneous information sources.

References

- Iñigo Alonso, Imanol Miranda, Eneko Agirre, and Mirella Lapata. 2025. [Tablet: A large-scale dataset for robust visual table understanding](#). *ArXiv*, abs/2509.21205.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xiong-Hui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Rongyao Fang, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, and 46 others. 2025a. [Qwen3-vl technical report](#).
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. [Qwen2.5-vl technical report](#). *ArXiv*, abs/2502.13923.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, SHIYANG LI, Xiyong Zhou, and William Yang Wang. 2019. [Tabfact: A large-scale dataset for table-based fact verification](#). *ArXiv*, abs/1909.02164.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matthew I. Beane, Ting-Hao 'Kenneth' Huang, Bryan R. Routledge, and William Yang Wang. 2021. [Finqa: A dataset of numerical reasoning over financial data](#). *ArXiv*, abs/2109.00122.
- Zhoujun Cheng, Haoyu Dong, Zhiruo Wang, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou, and Dongmei Zhang. 2021. [Hitab: A hierarchical table dataset for question answering and natural language generation](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Google DeepMind. 2026. [Gemma 4](#).
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Jun-Mei Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiaoling Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 179 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *ArXiv*, abs/2501.12948.
- Naihao Deng, Zhenjie Sun, Ruiqi He, Aman Sikka, Yulong Chen, Lin Ma, Yue Zhang, and Rada Mihalcea. 2024. [Tables as texts or images: Evaluating the table reasoning ability of llms and mllms](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Allen Institute for AI. 2025. [Molmo2: Open weights and data for vision-language models with video understanding and grounding](#).
- Jun-Peng Jiang, Yu Xia, Hai-Long Sun, Shiyin Lu, Qing-Guo Chen, Weihua Luo, Kaifu Zhang, De-Chuan Zhan, and Han-Jia Ye. 2025. [Multimodal tabular reasoning with privileged structured information](#). *ArXiv*, abs/2506.04088.
- Gemma Team Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ram' e, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean-Bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gael Liu, and 191 others. 2025. [Gemma 3 technical report](#). *ArXiv*, abs/2503.19786.
- Yoonsik Kim, Moonbin Yim, and Ka Yeon Song. 2024. [Tablevqa-bench: A visual question answering benchmark on multiple table domains](#). *ArXiv*, abs/2404.19205.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023. [Improved baselines with visual instruction tuning](#).
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. [Llava-next: Improved reasoning, ocr, and world knowledge](#).
- Suyash Vardhan Mathur, Jainit Sushil Bafna, Kunal Kartik, Harshita Khandelwal, Manish Shrivastava, Vivek Gupta, Mohit Bansal, and Dan Roth. 2024. [Knowledge-aware reasoning over multimodal semi-structured tables](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Meta. 2025. [Llama 4](#).
- Linyong Nan, Chia-Hsuan Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryscinski, Nick Schoelkopf, Riley Kong, Xiangru Tang, Murori Mutuma, Benjamin Rosand, Isabel Trindade, Renusree Bandaru, Jacob Cunningham, Caiming Xiong, and Dragomir R. Radev. 2021. [Fetaqa: Free-form table question answering](#). *Transactions of the Association for Computational Linguistics*, 10:35–49.
- OpenAI. 2026a. [Gpt-5.4](#).
- OpenAI. 2026b. [Gpt-5.4-mini](#).
- Panupong Pasupat and Percy Liang. 2015. [Compositional semantic parsing on semi-structured tables](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Qwen Team. 2026. [Qwen3.6-35B-A3B: Agentic coding power, now open to all](#).
- Qwen Team. 2026. [Qwen3.5: Accelerating productivity with native multimodal agents](#).
- Prasham Titiya, Jainil Trivedi, Chitta Baral, and Vivek Gupta. 2025. [Mmtbench: A unified benchmark for complex multimodal table reasoning](#). *ArXiv*, abs/2505.21771.

- Lanrui Wang, Mingyu Zheng, Hongyin Tang, Zheng Lin, Yanan Cao, Jingang Wang, Xunliang Cai, and Weiping Wang. 2025a. [Needleinatable: Exploring long-context capability of large language models towards long-structured tables](#). *ArXiv*, abs/2504.06560.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, and 44 others. 2025b. [InternV3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency](#). *ArXiv*, abs/2508.18265.
- Jian Wu, Linyi Yang, Dongyuan Li, Yuliang Ji, Manabu Okumura, and Yue Zhang. 2025. [Mmqa: Evaluating llms with multi-table multi-hop complex questions](#). In *International Conference on Learning Representations*.
- Bohao Yang, Yingji Zhang, Dong Liu, Andr e Freitas, and Chenghua Lin. 2025a. [Does table source matter? benchmarking and improving multimodal scientific table understanding and reasoning](#). *ArXiv*, abs/2501.13042.
- Zheyuan Yang, Lyuhao Chen, Arman Cohan, and Yilun Zhao. 2025b. [Table-r1: Inference-time scaling for table reasoning](#). *ArXiv*, abs/2505.23621.
- Tianyu Yu, Zefan Wang, Chongyi Wang, Fuwei Huang, Wenshuo Ma, Zhihui He, Tianchi Cai, Weize Chen, Yuxiang Huang, Yuanqian Zhao, Bokai Xu, Junbo Cui, Ying Xu, Liqing Ruan, Lu Zhang, Hanyu Liu, Jing Tang, Hongyuan Liu, Qi Guo, and 15 others. 2025. [Minicpm-v 4.5: Cooking efficient mllms via architecture, data, and training recipe](#). *ArXiv*, abs/2509.18154.
- Junwen Zhang, Pu Chen, and Yin Zhang. 2025. [Table-moe: Neuro-symbolic routing for structured expert reasoning in multimodal table understanding](#). *ArXiv*, abs/2506.21393.
- Tianshu Zhang, Xiang Yue, Yifei Li, and Huan Sun. 2023. [Tablellama: Towards open large generalist models for tables](#). *ArXiv*, abs/2311.09206.
- Weichao Zhao, Hao Feng, Qi Liu, Jingqun Tang, Shubo Wei, Binghong Wu, Lei Liao, Yongjie Ye, Hao Liu, Houqiang Li, and Can Huang. 2024a. [Tabpedia: Towards comprehensive visual table understanding with concept synergy](#). *ArXiv*, abs/2406.01326.
- Yilun Zhao, Yitao Long, Hongjun Liu, Ryo Kamoi, Linyong Nan, Lyuhao Chen, Yixin Liu, Xiangru Tang, Rui Zhang, and Arman Cohan. 2024b. [DocMath-eval: Evaluating math reasoning capabilities of LLMs in understanding long and specialized documents](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16103–16120, Bangkok, Thailand. Association for Computational Linguistics.
- Yilun Zhao, Zhenting Qi, Linyong Nan, Boyu Mi, Yixin Liu, Weijin Zou, Simeng Han, Ruizhe Chen, Xiangru Tang, Yumo Xu, Dragomir Radev, and Arman Cohan. 2023a. [QTSumm: Query-focused summarization over tabular data](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1157–1172, Singapore. Association for Computational Linguistics.
- Yilun Zhao, Chen Zhao, Linyong Nan, Zhenting Qi, Wenlin Zhang, Xiangru Tang, Boyu Mi, and Dragomir Radev. 2023b. [RobuT: A systematic study of table QA robustness against human-annotated adversarial perturbations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6064–6081, Toronto, Canada. Association for Computational Linguistics.
- Mingyu Zheng, Xinwei Feng, Qingyi Si, Qiaoqiao She, Zheng Lin, Wenbin Jiang, and Weiping Wang. 2024. [Multimodal table understanding](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Bangbang Zhou, Zuan Gao, Zixiao Wang, Boqiang Zhang, Yuxin Wang, Zhineng Chen, and Hongtao Xie. 2025a. [Syntab-llava: Enhancing multimodal table understanding with decoupled synthesis](#). *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24796–24806.
- Wei Zhou, Mohsen Mesgar, Heike Adel, and Annemarie Friedrich. 2025b. [Texts or images? a fine-grained analysis on the effectiveness of input representations and models for table question answering](#). *ArXiv*, abs/2505.14131.

ID	Education	Major	Assigned Task
1	PhD Student (Yr 3)	Computer Science	Attribute Labeling, QA Enhancement, Visual Rendering Audit, Quality Audit
2	Master (Yr 2)	Computer Science	Attribute Labeling, QA Enhancement
3	PhD Student (Yr 2)	Computer Science	Attribute Labeling, QA Enhancement, Visual Rendering Audit
4	PhD Student (Yr 4)	Computer Science	Attribute Labeling, QA Enhancement, Quality Audit
5	PhD Student (Yr 2)	Computer Science	Attribute Labeling, QA Enhancement, Visual Rendering Audit
6	Master (Yr 2)	Computer Science	Attribute Labeling, QA Enhancement
7	PhD Student (Yr 1)	Data Science	Attribute Labeling, QA Enhancement
8	Master (Yr 2)	Data Science	Attribute Labeling, QA Enhancement, Visual Rendering Audit
9	PhD Student (Yr 3)	Statistics	Attribute Labeling, QA Enhancement, Quality Audit
10	PhD Student (Yr 2)	Information Systems	Attribute Labeling, QA Enhancement
11	Master (Yr 2)	Economics	Attribute Labeling, QA Enhancement
12	PhD Student (Yr 1)	Mathematics	Attribute Labeling, QA Enhancement

Table 6: Profile of the 12 expert annotators recruited for TableVista. All annotators participated in *Attribute Labeling* (table layout, information richness, skill scores, and reasoning steps) and *QA Enhancement* (question rephrasing and deep synthesis) over the base set. Four annotators additionally conducted the *Visual Rendering Audit*, verifying answerability under structural and visual perturbations (e.g., *Missing* transformations). Three senior annotators further performed the final *Quality Audit*, reviewing the full benchmark for consistency, answerability, and difficulty calibration.

A Annotator Information

To ensure the high quality of TableVista, we recruited 12 expert annotators. All annotators are graduate students (Master’s or PhD) in quantitative or technical fields such as computer science, data science, statistics, and mathematics. Prior to the main annotation, each annotator completed a standardized 1-hour tutorial on our protocols and a 30-sample pilot annotation, which was reviewed and approved by the project leads before proceeding. Table 6 provides detailed profiles of the annotation team.

For quality assurance, every sample was reviewed by at least one annotator during the main annotation pass, and 10% of samples were independently cross-validated by a second annotator, with disagreements resolved through consensus discussion. Finally, all samples were passed through automated checks for answer uniqueness and format consistency before final approval.

B Quality Control Statistics

To provide full transparency regarding the rigor of our data construction pipeline, we report the rejection and revision statistics at every stage, based on our original annotation and audit logs. The pipeline proceeds in two phases—Base Set Construction and Visual Rendering Pipeline—each subject to multiple rounds of filtering, verification, and cross-

validation. Table 7 summarizes the quantitative outcomes of each stage.

Base Set Construction. We first aggregated raw records from the 14 public datasets and performed AI-assisted preliminary cleaning to remove low-quality samples while maintaining a balanced source distribution, retaining approximately 100,000 candidate samples. All 100,000 samples then underwent AI-assisted *Attribute Labeling* with manual verification as described in §3.1, covering table layout attributes and dimensions, *information richness* scores, *skill scores* across the four reasoning dimensions, and the number of *reasoning steps*. Based on these labeled attributes, we applied *Taxonomic Categorization and Multidimensional Filtering* to select 3,000 samples that jointly exhibit strong structural diversity and high reasoning complexity, corresponding to a 97.0% rejection rate. For the retained 3,000 samples, *QA Enhancement* was performed under the dual-track strategy described in §3.1: 2,264 samples were rephrased to increase linguistic diversity and reduce potential contamination, while the remaining 736 samples were authored from scratch as new multi-step reasoning questions grounded in the table and textual context.

Visual Rendering Pipeline. The 3,000 textual samples were rendered into the 10 visual settings

Phase	Stage	Initial	Removed / Revised	Rate
Base Set	Taxonomic Categorization and Multidimensional Filtering	100,000	97,000 removed	97.00%
	QA Enhancement – Question Rephrasing	3,000	2,264 rephrased	75.47%
	QA Enhancement – Deep Synthesis (regeneration)	3,000	736 regenerated	24.53%
Visual Rendering	Rendering Failure Handling	3,000	14 replaced	0.47%
	Answerability Verification & Calibration	30,000	912 regenerated	3.04%
	Cross Validation (10% of base samples, double-blind)	300	6 corrected	2.00%

Table 7: Rejection and revision statistics at each stage of the TableVista construction pipeline. The Base Set phase filters 100,000 attribute-labeled candidates down to 3,000 samples with high structural and reasoning diversity, followed by QA Enhancement in which every sample is either rephrased or rewritten from scratch. The Visual Rendering phase produces 30,000 images across 10 categories and applies three successive rounds of verification, yielding a final benchmark with fully verified answerability and minimal residual annotation disagreement.

described in §3.2, yielding 30,000 images in total. Three rounds of quality control were applied. First, in *Rendering Failure Handling*, 14 samples that failed rendering due to aspect ratio or context-length constraints were replaced to ensure completeness. Second, in *Answerability Verification and Parameter Calibration*, we repeatedly sampled 100 images per setting to assess post-render answerability and iteratively tuned the rendering parameters until the pipeline reached 97.2% answerability; after fixing the parameters, we manually verified answerability for all 30,000 rendered images and regenerated 912 images (3.04%) whose content had been compromised by the visual transformations. Finally, as part of the *human-in-the-loop* audit introduced in §3.3, 10% of the base samples (300 samples together with all their rendered variants) were subjected to double-blind cross-validation by two independent annotators; only 6 samples required minor corrections, corresponding to a 2.0% disagreement rate and indicating that the earlier audit stages had already resolved the vast majority of quality issues.

Overall, these statistics demonstrate that TableVista is the product of a multi-stage, human-in-the-loop pipeline in which every sample is filtered, enhanced, rendered, and independently verified.

C Experiment Setup

C.1 Configuration of Evaluated Models

Table 8 lists the full inference configuration of each evaluated model. Rather than imposing a single uniform decoding setting, we follow the sampling parameters (temperature, top- p , top- k , and repetition penalty) recommended by each model’s offi-

cial release, since recommended defaults vary substantially across model families; models without an explicit sampling recommendation are evaluated with greedy decoding. The maximum number of generated tokens is set to 8,192 for open-source models, except for the LLaVA family, whose native 4,096-token context window restricts generation to 2,048 tokens. For closed-source models accessed through official API endpoints, we use the provider’s default sampling configuration and cap the completion length at 16,000 tokens. All open-source models are served locally via vLLM with tensor-parallel sizes ranging from 1 to 8 depending on model scale. Our experiments are conducted on a workstation equipped with multiple NVIDIA H20 GPUs.

C.2 Prompts

We release the full set of prompts used in our evaluation. As discussed in §4.4, all main results in §4.2 and §4.3 adopt the *direct-output* template (Figure 4), while the CoT variant (Figure 5) is used in the CoT-vs-direct-output analysis. Each template comes with two modalities: a *general* version for the standard image+text setting, and a *vision-only* version for the *Screen Capture* and *Simulated Photo* configurations where the question is rendered into the image. In addition, the LLM-as-Judge prompt (Figure 6) is applied as a secondary verification layer on predictions that fail the Exact Match criterion.

Organization	Model	Version	Serving Backend	Temp.	Max Tokens
<i>Proprietary Models</i>					
OpenAI	GPT-5.4	gpt-5.4-0305-global	API	–	16,000
	GPT-5.4-mini	gpt-5.4-mini-0317-global	API	–	16,000
<i>Open-Source Foundation Models</i>					
Qwen	Qwen3.6-35B-A3B	Qwen/Qwen3.6-35B-A3B	vLLM	1.0	8,192
	Qwen3.5-122B-A10B	Qwen/Qwen3.5-122B-A10B	vLLM	1.0	8,192
	Qwen3.5-35B-A3B	Qwen/Qwen3.5-35B-A3B	vLLM	1.0	8,192
	Qwen3.5-27B	Qwen/Qwen3.5-27B	vLLM	1.0	8,192
	Qwen3.5-9B	Qwen/Qwen3.5-9B	vLLM	1.0	8,192
	Qwen3-VL-30B-A3B	Qwen/Qwen3-VL-30B-A3B-Instruct	vLLM	0.7	8,192
	Qwen3-VL-8B	Qwen/Qwen3-VL-8B-Instruct	vLLM	0.7	8,192
	Qwen2.5-VL-72B	Qwen/Qwen2.5-VL-72B-Instruct	vLLM	0.0	8,192
	Qwen2.5-VL-32B	Qwen/Qwen2.5-VL-32B-Instruct	vLLM	0.0	8,192
InternVL	Qwen2.5-VL-7B	Qwen/Qwen2.5-VL-7B-Instruct	vLLM	0.0	8,192
	InternVL3.5-30B-A3B	OpenGVLab/InternVL3_5-30B-A3B	vLLM	0.0	8,192
	InternVL3.5-14B	OpenGVLab/InternVL3_5-14B	vLLM	0.0	8,192
Meta	InternVL3.5-8B	OpenGVLab/InternVL3_5-8B	vLLM	0.0	8,192
	Llama-4-Maverick	meta-llama/Llama-4-Maverick-17B-128E-Instruct-FP8	vLLM	0.6	8,192
Google	Llama-4-Scout	meta-llama/Llama-4-Scout-17B-16E-Instruct	vLLM	0.6	8,192
	Gemma-4-31B-it	google/gemma-4-31b-it	vLLM	1.0	8,192
	Gemma-4-26B-A4B-it	google/gemma-4-26B-A4B-it	vLLM	1.0	8,192
	Gemma-4-E4B-it	google/gemma-4-E4B-it	vLLM	1.0	8,192
	Gemma-3-27B-it	google/gemma-3-27b-it	vLLM	1.0	8,192
Allen AI	Gemma-3-12B-it	google/gemma-3-12b-it	vLLM	1.0	8,192
	Molmo2-8B	allenai/Molmo2-8B	vLLM	0.0	8,192
OpenBMB	MiniCPM-V-4.5	openbmb/MiniCPM-V-4_5	vLLM	0.6	8,192
LLaVA	LLaVA-v1.6-Vicuna-13B	llava-hf/llava-v1.6-vicuna-13b-hf	vLLM	0.0	2,048
	LLaVA-v1.6-Vicuna-7B	llava-hf/llava-v1.6-vicuna-7b-hf	vLLM	0.0	2,048
	LLaVA-v1.5-7B	llava-hf/llava-1.5-7b-hf	vLLM	0.0	2,048
Table-LLaVA	Table-LLaVA-v1.5-13B	SpursgoZmy/table-llava-v1.5-13b-hf	vLLM	0.0	2,048
	Table-LLaVA-v1.5-7B	SpursgoZmy/table-llava-v1.5-7b-hf	vLLM	0.0	2,048

Table 8: Details of the multimodal foundation models evaluated in TableVista and their inference configurations. The **Serving Backend** indicates how the model was accessed in our experiments (API for cloud endpoints; vLLM for local deployment on NVIDIA H20 GPUs). **Temperature** follows each model’s official recommendation; “–” indicates that we use the provider’s default sampling configuration for API models. Other sampling parameters (top- p , top- k , and repetition penalty, when applicable) also follow each model’s official recommendation. **Max Tokens** is the upper bound on generated tokens per response.

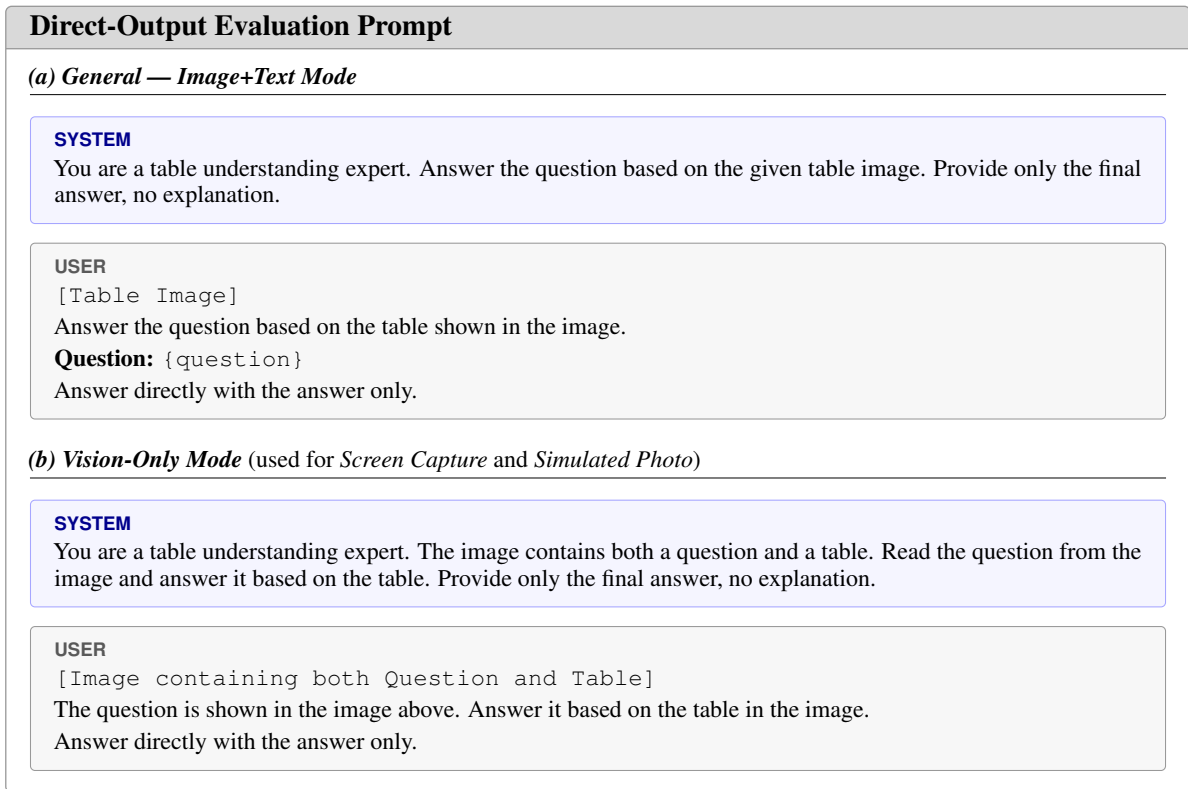


Figure 4: Direct-output evaluation prompt. (a) The general template supplies the table image and the textual question as separate inputs. (b) The vision-only template is used when the question is rendered into the image together with the table, requiring the model to jointly parse both from a single visual input.

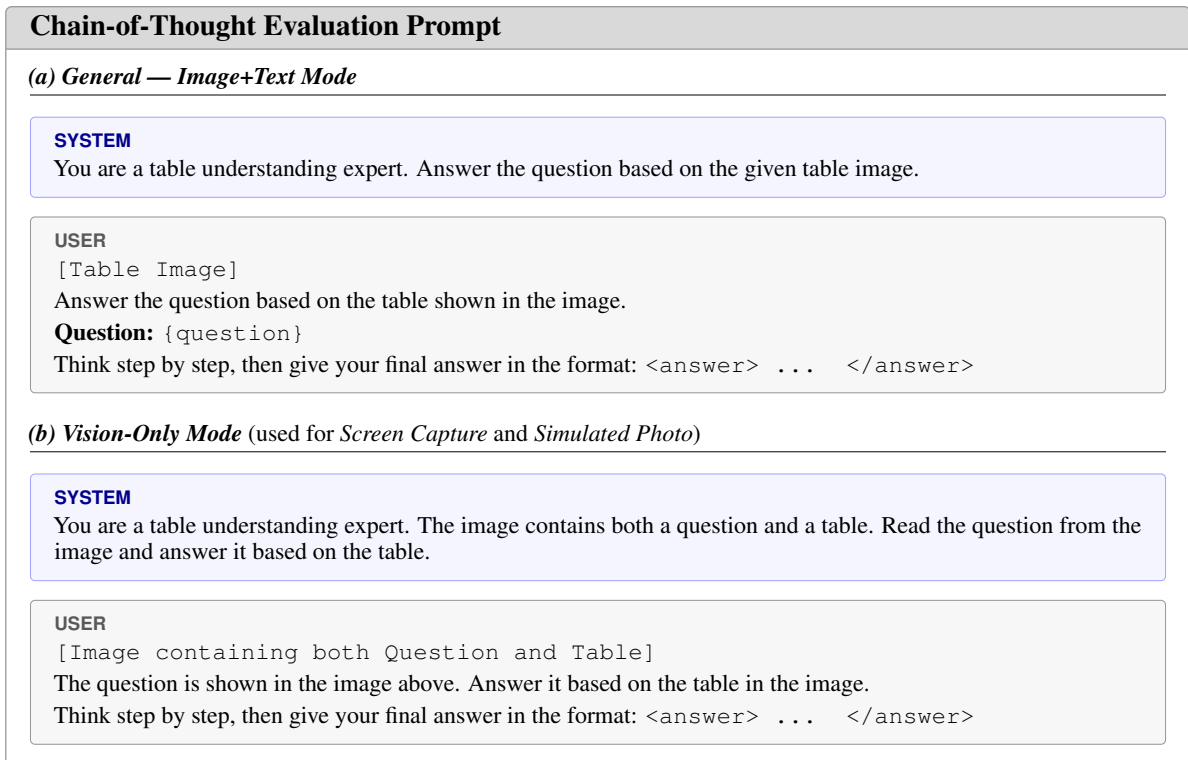


Figure 5: Chain-of-Thought evaluation prompt. (a) The general template elicits step-by-step reasoning before the final tagged answer. (b) The vision-only template applies the same CoT protocol when the question and table are provided together as a single image. For models supporting native thinking mode (e.g., Qwen3-VL, Qwen3.5), we additionally enable `enable_thinking=True` at inference.

```
LLM-as-Judge Evaluation Prompt
Applied to predictions that fail Exact Match, using GPT-5-mini as the judge model.
USER
You are evaluating a table question answering system.
Question: {question}
Ground truth answer: {answer}
Predicted answer: {predicted}
Mark as CORRECT ('yes') if any of the following apply:
• The prediction matches the gold answer exactly or semantically (case, abbreviation, minor formatting differences are OK).
• The prediction contains the gold answer and only adds valid supplementary information such as units or clarifying context (e.g., gold='713', pred='713 thousand shares' is OK).
Mark as INCORRECT ('no') otherwise.
Numeric substring trap: gold='2' ≠ pred='20', gold='13' ≠ pred='13.5'.
Reply with ONLY 'yes' or 'no'.
```

Figure 6: LLM-as-Judge evaluation prompt. Applied as a secondary verification layer on predictions that fail the Exact Match criterion, tolerating formatting variations while rejecting numeric substring collisions. We use GPT-5-mini as the judge throughout all experiments.

D Qualitative Analysis

To gain deeper insights into the failure modes of multimodal foundation models beyond aggregate metrics, we present detailed case studies in TableVista. We examine the impact of visual perturbations on reasoning consistency in [Figure 7](#) and the challenges of spatial grounding in complex layouts in [Figure 8](#).

E Visual Examples

To demonstrate the structural diversity and complexity of the TableVista benchmark, we present representative samples across different categories. [Figure 9](#) displays examples of the five table structural types defined in our taxonomy, ranging from standard grids to hierarchical and multi-table layouts, verifying the breadth of challenges covered in our dataset.

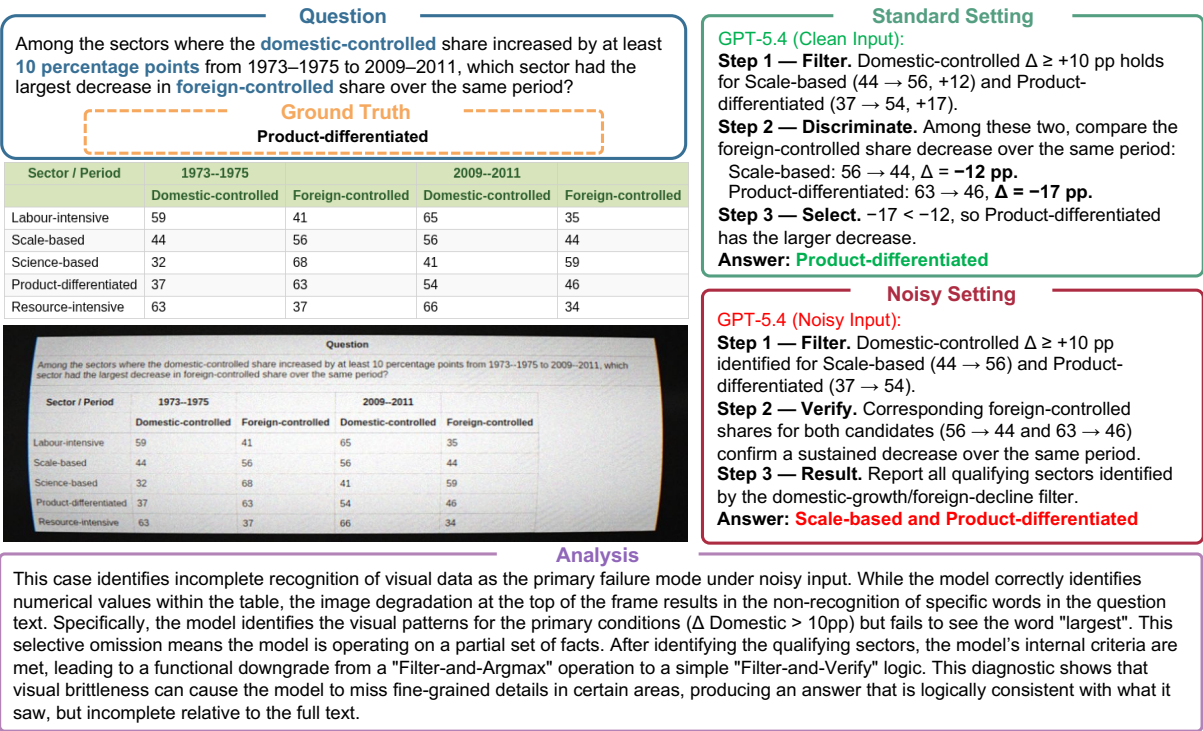


Figure 7: Qualitative analysis of visual brittleness in TableVista.

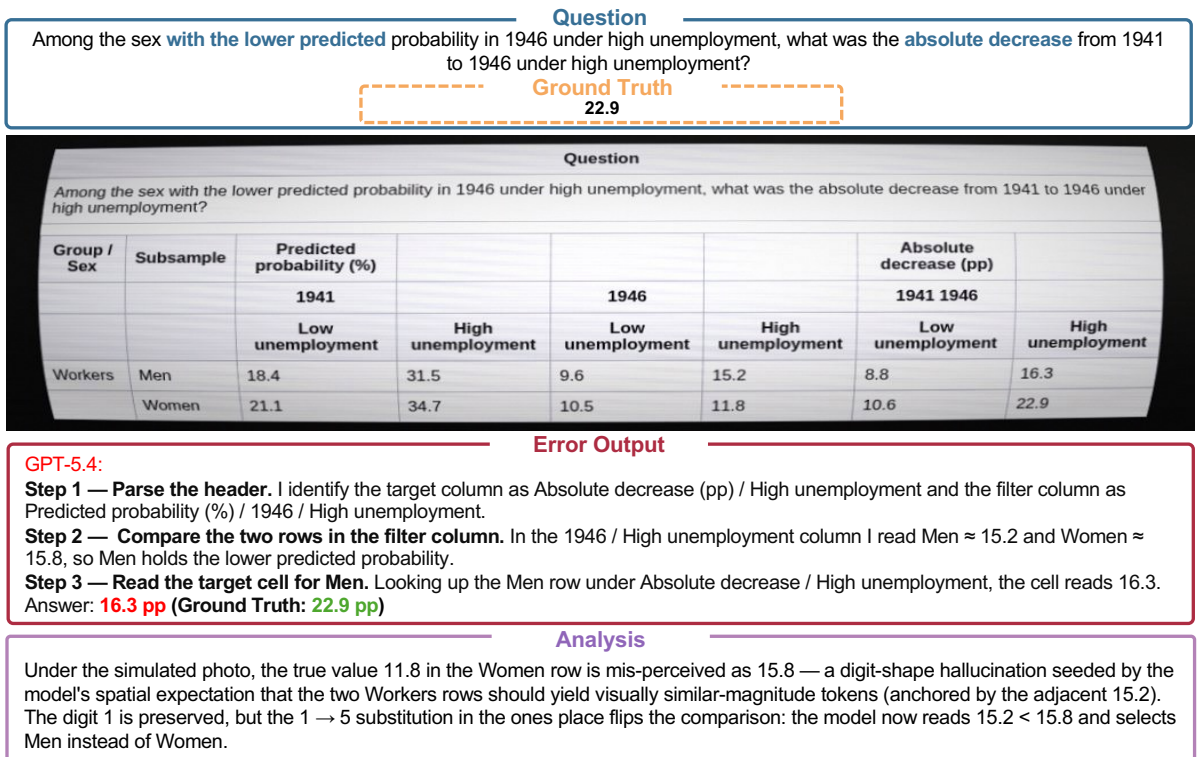


Figure 8: Qualitative analysis of spatial alignment failures in TableVista.

Context
 Performance graph. The following table compares the five-year cumulative return of our common shares, the S&P 500 Index, and the NAREIT All Equity Index, a peer group index. The figures assume that \$100 was invested on December 31, 2009 in each series and that all dividends were reinvested without the payment of any commissions.

Type	Entity / Year	Base		Mid-period		End-period	
		2009	2010	2011	2012	2013	2014
Company	vornado realty trust	\$100	\$123	\$118	\$128	\$147	\$201
Benchmark	s&p 500 index	100	115	117	136	180	205
	the nareit all equity index	100	128	139	166	171	218

(a) Text-Mixed Table (Web)

Code / Meaning	Color Description
9	red
5	green
1	yellow
4	blue
7	black
2	white
8	purple
3	gray

Rank / Nation	Team	Medal tally				Total
		Gold	Silver	Bronze		
1	soviet union (urs)	11	9	9	29	
2	east germany (gdr)	9	10	6	25	
3	switzerland (sui)	5	5	5	15	
4	finland (fin)	4	1	2	7	
5	sweden (swe)	4	0	2	6	
6	austria (aut)	3	5	2	10	
7	netherlands (ned)	3	2	2	7	
8	west germany (frg)	2	4	2	8	
9	united states (usa)	2	1	3	6	
10	italy (ita)	2	1	2	5	
13	canada (can)	0	2	3	5	

(b) Simple Table (Custom)

Identifiers		Product Info		Pricing		Descriptions	
Product ID	Color Code	Category	Product Name	Typical Buying	Typical Selling	Product Description	Other Details
			Price	Price			
1	4	Spices	cumin	2878.3		et	None
2	2	Spices	peper	352447.2874677	1892070.2803543	rerum	None
3	9	Herbs	basil	503.8431967	0.1859512	officia	None
4	1	Herbs	borage	10377614.847385		blandititis	None
5	4	Spices	chilli		39446	eius	None
6	4	Seeds	ginger	5.578	52735.6101	doloribus	None
7	9	Seeds	sesame	1284268.0659	68205825.7	et	None
8	9	Herbs	caraway	24493		nulla	None
9	2	Herbs	catnip	12008702.623	21577.891642	vel	None
10	5	Seeds	flax	339404395.7	59622629.74	et	None
11	7	Herbs	chervil			minus	None
12	4	Seeds	voluptatem	162	149	officia	None
13	5	Spices	cinnam	1686539.4	17595111.4	nisi	None
14	4	Seeds	lotus	43221310.465574	63589.4054376	exercitationem	None
15	2	Herbs	laurel		57857	ut	None

(c) Multi Table (LaTeX)

Year / Quarter	Half	Revisions to quarterly GDP growth (percentage points)				
		Output approach		Expenditure approach		
		GDP at market prices		GVA at basic prices	GDP at market prices	Domestic demand
2013	H1	Q1: 0.0	0.1	Q1: 0.1	0.2	
		Q2: 0.1	0.0	Q2: 0.0	0.1	
	H2	Q3: 0.2	0.1	Q3: 0.2	0.1	
		Q4: 0.1	0.2	Q4: 0.1	0.0	
2014	H1	Q1: 0.2	0.1	Q1: 0.3	0.2	
		Q2: 0.1	0.1	Q2: 0.2	0.1	
	H2	Q3: 0.3	0.2	Q3: 0.4	0.3	
		Q4: 0.2	0.2	Q4: 0.3	0.2	
2015	H1	Q1: 0.0	-0.1	Q1: -0.1	0.0	
		Q2: 0.2	0.1	Q2: 0.2	0.1	
	H2	Q3: 0.1	0.0	Q3: 0.1	0.1	
		Q4: 0.0	0.1	Q4: 0.0	-0.1	

(d) Complex Table (Custom)

1996 Italian motorcycle Grand Prix				
Entry details				
Pos	Rider	Manufacturer	Classification	Points
Classified finishers				
Points-paying positions (1-15)				
1	Max Biaggi	Aprilia	40:36.299	25
2	Marcelino Lucchi	Aprilia	+6.914	20
3	Ralf Waldmann	Honda	+18.294	16
4	Olivier Jacque	Honda	+18.607	13
5	Jurgen Fuchs	Honda	+31.617	11
6	Tetsuya Harada	Yamaha	+31.649	10
7	Jean-Philippe Ruggia	Honda	+31.683	9
8	Nobuatsu Aoki	Honda	+47.701	8
9	Luis d'Antin	Honda	+51.186	7
10	Cristiano Migliorati	Honda	+51.556	6
11	Massimo Ottobro	Aprilia	+51.563	5
12	Eskil Suter	Aprilia	+51.613	4
13	Luca Boscoscuro	Aprilia	+51.617	3
14	Regis Laconi	Honda	+52.294	2
15	Jurgen vd Goorbergh	Honda	+52.569	1
Non-scoring classified finishers (16-28)				
16	Sebastian Porto	Aprilia	+1:03.619	
17	Jamie Robinson	Aprilia	+1:06.917	
18	Olivier Petrucci	Aprilia	+1:06.945	
19	Alessandro Antonello	Aprilia	+1:07.099	
20	Osamu Miyazaki	Aprilia	+1:32.873	
21	Davide Bulega	Aprilia	+1:32.902	
22	Yasumasa Hatakeyama	Honda	+1:34.039	
23	Franco Battaini	Aprilia	+1:34.363	
24	Sete Gibernau	Honda	+1:34.931	
25	Christian Boudinot	Aprilia	+1:41.263	
26	José Barresi	Yamaha	+1 Lap	
27	Roberto Rolfo	Aprilia	+1 Lap	
28	Cristophe Cogan	Honda	+1 Lap	
Retirements				
Ret	Gianluigi Scalvini	Honda	Retirement	
Ret	Roberto Locatelli	Aprilia	Retirement	
Ret	Takeshi Tsujimura	Honda	Retirement	

(e) Long Table (Excel)

Figure 9: Visual examples for different table structures.