

VisDoT : Enhancing Visual Reasoning through Human-Like Interpretation Grounding and Decomposition of Thought

Eunsoo Lee¹, Jeongwoo Lee², Minki Hong¹, Jangho Choi¹, Jihie Kim^{1†}

¹Department of Computer Science and Artificial Intelligence, Dongguk University

²Department of Electronics and Electrical Engineering, Dongguk University

{dmstn7432, jwlee0519, jackyh1, 2025120382}@dgu.ac.kr, jihie.kim@dgu.edu

Abstract

Large vision-language models (LVLMs) struggle to reliably detect visual primitives in charts and align them with semantic representations, which severely limits their performance on complex visual reasoning. This lack of perceptual grounding constitutes a major bottleneck for chart-based reasoning. We propose VisDoT, a framework that enhances visual reasoning through human-like interpretation grounding. We formalize four perceptual tasks based on the theory of graphical perception such as position and length. Building on this foundation, we introduce decomposition-of-thought (DoT) prompting, which sequentially separates questions into visual perception sub-questions and logic sub-questions. Fine-tuning InternVL with VisDoT achieves a +11.2% improvement on ChartQA and surpasses GPT-4o on the more challenging ChartQAPro benchmark. On the newly introduced VisDoTQA benchmark, the model improves by +33.2%. Furthermore, consistent zero-shot gains on diverse open-domain VQA benchmarks confirm the generalizability of the perception-logic separation strategy for visual question answering in general. VisDoT leverages human-like perception to enhance visual grounding, achieving state-of-the-art chart understanding and interpretable visual reasoning.

1 Introduction

One of the fundamental challenges for vision-language models (VLMs) is their limited ability to interpret Visualized Data such as charts, graphs, infographics, and dashboards and to perform reasoning over interrelated visual elements. In real-world scenarios, large vision-language models (LVLMs) often fail to reliably detect visual primitives and align them with semantic representations when users issue queries or instructions about visual data without explicitly mentioning identifiers such

Q: How much higher is the “Higher” response from the leftmost political group compared to the group with the lowest such response?




Figure 1: Comparison of visual reasoning between existing LVLMs and VisDoT. While LVLMs fail to visual perception in spatial structure, VisDoT leverages decomposition-of-thought (DoT) to accurately infer answers through sequential visual analysis.

as legend labels or axis names. This results in substantial performance degradation (Figure 1). Prior work on Visualized Data understanding has focused primarily on adapting VLMs through instruction tuning and chain-of-thought (CoT) supervision (Meng et al., 2024; Han et al., 2023). These approaches are typically limited to simple text-based keyword-value mappings and lack robust grounding mechanisms that align visual primitives such as color, spatial coordinates, and shape with their corresponding real-world entities. As a result, their performance degrades substantially on tasks that require high-level perceptual alignment, such as legend identification and multi-object comparison (Masry et al., 2022, 2025a). Similarly, while CoT supervision has shown effectiveness for text-only reasoning, it provides only limited gains in visual reasoning

contexts that require logical reasoning (Cărbune et al., 2024; Wu et al., 2025).

In contrast, recent studies have demonstrated the potential of LVLMs for visual question answering (VQA) by integrating visual elements into instruction-following datasets (Masry et al., 2025b). In parallel, research on modular VLMs has shown that disentangling reasoning, visual understanding, and language understanding is an effective strategy for VQA (Amizadeh et al., 2020). Furthermore, recent work suggests that a similar separation between visual perception and logical reasoning is also necessary for LVLMs (Zhang et al., 2024b).

Based on these insights, we pose two central research questions: (1) *How can LVLMs be effectively adapted for visualized data-based question answering?*, and (2) *Can task decomposition strategies be effectively extended to LVLMs?*

To address these, we propose VisDoT, a framework that enhances the grounding and reasoning capabilities of LVLMs. VisDoT leverages graphical perception theory to improve visual grounding and employs systematic task decomposition to enable structured reasoning. Building on the theory of human graphical perception (Cleveland and McGill, 1984), we formalize core perceptual tasks: Position, Length, Pattern, and Extract. These tasks serve as a foundation for aligning model attention with human perceptual principles. We further introduce a decomposition-based prompting strategy, decomposition-of-thought (DoT), which separates questions into perceptual and logical reasoning stages to enable sequential problem solving. The LVLM is trained to perform the four perceptual tasks and to adopt the DoT reasoning strategy.

A model trained with the proposed framework provides two key benefits: (1) improved robustness to complex and underspecified visual queries, and (2) enhanced interpretability of VLM failure cases. Experiments on ChartQA (Masry et al., 2022), ChartQAPro (Masry et al., 2025a), and cross-domain visual reasoning benchmarks (Li et al., 2023; Yue et al., 2024) yield state-of-the-art performance, underscoring the framework’s effectiveness for real-world visual reasoning. In particular, the model achieves an 11.2% gain on ChartQA and matches or surpasses GPT-4o (OpenAI, 2024) on the more challenging ChartQAPro and our VisDoTQA benchmark. Notably, applied to the open-domain VQA benchmarks POPE (Li et al., 2023) and MMMU (Yue et al., 2024), our DoT prompting raises performance by 1.43% and

2.2% absolute over an identical CoT-trained backbone under the same compute budget. By requiring each sub-question to first localise the relevant visual element before reasoning, DoT removes the grounding errors that limit CoT, showing that the approach generalises well beyond chart tasks.

Our contributions are summarized as follows:

- We formalize four core perceptual tasks grounded in graphical perception theory, establishing the foundation for modeling LVLMs that emulate human decoding of data visualizations.
- We introduce a novel DoT strategy that enables LVLMs to emulate human-like visual interpretation by separating complex questions into perceptual and logical steps.
- We construct a perception-following dataset that combines these perceptual tasks with DoT-based prompting, enabling chart understanding, visual-linguistic grounding, and compositional reasoning.
- Our trained models achieve state-of-the-art performance across multiple visual reasoning benchmarks including VQA tasks, demonstrating the effectiveness of the VisDoT framework in grounding visual-language inference.

2 Related Work

2.1 Reasoning from Visualized Data

Recent advances in LVLMs have spurred a growing body of research aimed at improving structured visual reasoning for complex inputs such as charts. Pix2Struct (Lee et al., 2023) improves layout understanding through markup-based reconstruction pretraining, while Matcha (Liu et al., 2024) focuses on chart de-rendering and numerical reasoning. Alongside the widespread adoption of instruction tuning, several studies have proposed adapting LVLMs to visualized data understanding by leveraging synthetic or large-scale instruction datasets (Han et al., 2023; Masry et al., 2024, 2025b). More recently, ChartGemma (Masry et al., 2025b) attempts to enhance the visual capabilities of LVLMs by incorporating a visual-element-based instruction tuning dataset into the training data. Chart-based Reasoning (Cărbune et al., 2024) reports improved VQA performance by generating intermediate reasoning processes or structured evidence. Similarly, prior work (Han




Figure 2: An overview of our framework.

et al., 2023; Masry et al., 2024, 2025b) improves chart reasoning performance by augmenting training datasets with CoT annotations. However, recent findings (Zhang et al., 2024b) emphasize that CoT (Wei et al., 2022) and Self-Consistency (Wang et al., 2022) have limited effectiveness in vision tasks.

2.2 Decomposition Strategies for Complex Visual Reasoning

To address complex visual reasoning, prior work has proposed various decomposition strategies, including modular execution and question decomposition. As a representative early approach, NS-VQA (Amizadeh et al., 2020) resolves this challenge by modularizing recognition and inference through a neuro-symbolic architecture. More recently, (Wang et al., 2023; Hu et al., 2023) improve VQA performance by invoking executable programs and tools or leveraging multiple agents. However, iterative modular execution incurs high computational costs and makes it difficult to transfer the approach to a single end-to-end sVLM. Question decomposition methods (Khan et al., 2023; Zhang et al., 2024a) suggest that splitting complex queries into sub-questions can improve LVLM performance. However, (Khan et al., 2023) does not handle a large number of sub-questions effectively, and (Zhang et al., 2024a) mainly focuses on demonstrating the potential of question decomposition strategies in LVLMs. In contrast, VisDoT integrates question decomposition and reasoning within a single system and a single reasoning process. This enables scalable annotation generation and effective knowledge transfer to compact LVLMs.

Task	Description
Position	Compares object positions along a common scale (e.g., x- or y-axis) to determine relative order. As the most accurate perceptual channel for conveying quantitative information, it plays a critical role in effective visual communication.
Length	A distortion-free visual attribute across visualized data types (Stevens, 2017), used as a secondary cue to position.
Pattern	Links pattern cues to legends and data to distinguish categories. Assesses visual label mapping ability.
Extract	Reads explicitly shown values. Evaluates numerical recognition similar to standard QA tasks.

Table 1: Core perceptual tasks derived from graphical perception theory, designed to align LVLM attention with structured visual features for visualized data reasoning. A detailed description of each task is provided in Appendix A.1

3 Approach

We propose VisDoT, a framework that enhances LVLMs’ visual reasoning by emulating human visual decoding and incorporating LVLM-specialized inference. As illustrated in Figure 2, VisDoT comprises two components: (1) a perception-following dataset generation grounded in human perceptual priors, and (2) a DoT strategy that splits complex queries into sequential perception and logic sub-questions. This section introduces the motivation and architecture of VisDoT, highlighting how each component contributes to complex visual reasoning. Section 3.1 describes the construction of perception-following questions based on structured visual representations. Section 3.2 presents the DoT strategy, which splits queries into perception–logic stages and produces aligned answers.




Figure 3: Perception-following QA examples for the four task types in Table 1. Each question is decomposed into perception and logic sub-questions using the DoT prompt, enabling structured and interpretable chart reasoning.

3.1 Perception-following Question for LVLMS

Human interpretation of visualized data involves decoding at the level of perceptual units, grounded in visual features such as position, length, or angle. Cognitive psychology research has shown that this process is hierarchically organized according to the accuracy and perceptual salience of these visual attributes (Cleveland and McGill, 1984). However, existing LVM training datasets predominantly consist of instruction-level QA pairs, making it difficult for models to learn or emulate such perceptual decoding processes (Vogel et al., 2025).

Perceptual Task. To lay the groundwork for modeling LVLMS that emulate human cognitive structures, we define a set of perceptual tasks (position, length, pattern, extract) grounded in human graphical cognition theory (Cleveland and McGill, 1984). As shown in Table 1, these tasks provide the basis for constructing training data that guides model attention to align with human visual information processing.

Perception-following Question Generation. Building on the defined perceptual tasks (Table 1), we develop an automated question generation algorithm (red arrow in Figure 2) that constructs questions requiring visual perception-based reasoning. Depending on task characteristics, the

generated questions may involve multi-object and multi-step inference. As illustrated in Figure 3, they range from OCR-style queries (e.g., identifying "major" or "not", corresponding to sub-figure (d)) to those with explicit visual descriptors (e.g., "dark red" or "light gray", as shown in sub-figure (c)), and also include questions that require implicit visual interpretation based on chart axes or guide lines (sub-figures (a) and (b)). This approach enables the generation of questions that evaluate the model's ability to interpret charts, align visual elements with linguistic references, and perform grounded reasoning. The full prompting template used for perception-following question generation is detailed in Appendix D.1.

3.2 Decomposition-of-Thought Method

Modular approaches that explicitly separate visual perception from logical reasoning have demonstrated strong performance and interpretability in complex VQA tasks, such as NS-VQA. In contrast, LVLMS typically employ language-centric reasoning strategies in multi-modal settings such as CoT. While these strategies are effective for symbolic tasks, they have shown limited effectiveness in scenarios that require grounded visual reasoning, as empirically demonstrated in Section 4.5.

Motivated by this gap, we extend the principle




Figure 4: The model is guided to decompose a complex visual question into perception and logic sub-questions (Question Decomposition) and generate intermediate reasoning steps sequentially (Problem Solving), enabling structured and interpretable visual inference.

of separating visual perception and logical reasoning to LVLMs and develop an LVLM-specialized visual reasoning strategy. In doing so, we redefine VQA as a compositional task and introduce the DoT method to guide LVLMs in solving these tasks. Finally, we generate answers by accumulating the phase-wise reasoning process enabled by the VisDoT framework as illustrated by the [blue arrow](#) in Figure 2.

Visual Question Answering definition. Standard LVLM-based VQA aims to maximize the conditional likelihood of the answer A given an image I and a question Q , which can be formalized as the following objective (Khan et al., 2023):

$$\arg \max_A P(A | I, Q) \quad (1)$$

Following the separating visual perception and logical reasoning principle, we redefine VQA by formulating a question Q as a compositional query that integrates visual perception and logical reasoning.

$$P(A | I, Q) = \sum_{\{Q_1, \dots, Q_n\}} P(\{Q_1^p, \dots, Q_n^l\} | Q) \cdot \prod_{i=1}^n P(A_i | I, Q_i, A_{<i})$$

The compositional question Q is represented as a set of sub-questions Q_i , each designed to solve a specific sub-task. Each sub-question Q_i is classified as either a perception-oriented sub-question

Q_i^p or a logic-oriented sub-question Q_i^l . We do not impose constraints on the number of perception or logic-oriented sub-questions, which allows flexible decomposition of complex tasks into arbitrarily many sub-questions. For each sub-question Q_i , the model generates an intermediate answer A_i sequentially. The final answer A_n is obtained by accumulating these intermediate answers. Each intermediate answer A_i is generated based on the image I and the preceding context $A_{<i}$, enabling context-aware multi-step reasoning. The first term $P(\{Q_1^p, Q_2^p, Q_3^l, \dots, Q_n^l\} | Q)$ represents the probability of decomposing the complex question Q into an ordered sequence of perceptual Q^p and logical Q^l sub-questions (Zhang et al., 2024b). We assume that decompositions that prioritize perception first-order are generally more plausible (Liu et al., 2025). The second term, $\prod_{i=1}^n P(A_i | I, Q_i, A_{<i})$, models the likelihood of accurately answering each sub-question based on image-based and context-aware reasoning. Thus, the most reliable answers are obtained when the question is logically decomposed and each step is reasoned correctly.

DoT Prompt. To effectively perform chart reasoning as defined in Equation 2, we design a DoT prompt that explicitly guides LVLMs in structured answer generation. As illustrated in Figure 4, the DoT prompt organizes the reasoning process into two phases.

Phase 1: Question Decomposition. The model decomposes a compositional question Q into an ordered set of sub-questions consisting of perception-oriented Q_i^p and logic-oriented Q_i^l queries. The prompt explicitly enforces that perception-oriented sub-questions Q_i^p are generated before logic-oriented Q_i^l ones. As illustrated in Figure 3(c), the prompt is designed to explicitly prioritize perception-oriented sub-questions (e.g., retrieving individual values or identifying visual properties) before logic-oriented ones (e.g., computing differences or performing comparisons). For example, questions 1 and 2 guide the model to extract visual information regarding chart components, while question 3 introduces a logical operation based on the previously grounded values.

Phase 2: Problem Solving. The model sequentially answers each sub-question Q_i while explicitly accumulating intermediate reasoning steps. The resulting output (yellow box in Figure 4) satisfies both components of Equation 2, thereby enabling structured and interpretable reasoning. The

Task Type	# QA Pairs
Position	66,617
Length	31,614
Pattern	61,268
Extract	174,411
Total	331,969

Table 2: Distribution of QA pairs across the four perceptual reasoning tasks in the VisDoTQA dataset.

full prompt used in our framework is provided in Appendix D. To verify that the proposed DoT prompting is not limited to chart-centric scenarios, we further evaluate its generalisation in Section 4.5 on two open-domain VQA benchmarks that do *not* involve charts. The same DoT template yields over CoT while markedly reducing visual-grounding errors, confirming that the perception-first decomposition strategy extends beyond chart reasoning to generic VQA tasks.

DoT-based Perception-following Answer Generation. As illustrated by the [blue arrow](#) in Figure 2, the LVLM generates answers that include the DoT reasoning process guided by the DoT prompt. Figure 3 shows that the model produces reasoning trajectories tailored to the given image and question. As seen in Figure 3 (b), the model generates sub-questions that sequentially query multiple objects and constructs stepwise logic-oriented questions to handle complex reasoning. The answers generated in this manner explicitly reveal the fine-grained reasoning process, which enhances the interpretability of reasoning paths and improves adaptability to complex visual queries. Finally, we train the LVLM on VisDoTQA, which is constructed using VisDoT, and conduct systematic evaluations against existing visual reasoning approaches on multiple benchmarks. The example instances from the generated dataset are shown in Appendix E.

4 Experiment

4.1 Experimental Setup

Data. We developed VisDoTQA using the VisDoT framework based on 16,167 chart images collected from Pew Research ([Pew Research Center, 2024](#)), Statista ([Statista, 2024](#)), Our World in Data (OWID) ([Our World in Data, 2024](#)), and the OECD ([OECD, 2024](#)). The distribution of the dataset is presented in Table 2. We fine-tune open-source models on VisDoTQA to validate the effectiveness of the VisDoT framework. Question gener-

ation in Figure 2 ([red arrow](#)) is performed using the closed-source GPT-4o ([OpenAI, 2024](#)), while answer generation ([blue arrow](#)) is conducted with the open-source LLaMA-3.2-90B ([Dubey et al., 2024](#)).

VisDoTQA test set. To prevent evaluation leakage, we hold out 609 charts that are never used in training or prompt development to construct VisDoTQA test set. Questions are generated with GPT-4o ([OpenAI, 2024](#)). For answer labeling, we apply the DoT prompt to both **Llama-3.2-90B-Vision** ([Dubey et al., 2024](#)) and **GPT-4o** ([OpenAI, 2024](#)); if the two models produce identical final answers, we adopt it as the ground truth, otherwise a human verifies and edits the label.

Baseline. (i) As model baselines, we compare InternVL2.5 ([Chen et al., 2024](#)), Qwen2.5 ([Bai et al., 2025](#)), and Gemma 3 ([Team et al., 2025](#)), a chart-specialized SOTA model ChartGemma ([Masry et al., 2025b](#)), and closed-source models GPT-4o ([OpenAI, 2024](#)), Gemini-Flash 2.0 ([Comanici et al., 2025](#)). (ii) As methodology baselines on the same student backbone, we compare two chart distillation/data-synthesis pipelines with publicly released instruction data (ChartGemma ([Masry et al., 2025b](#)) and ECD ([Yang et al., 2025](#))). Full experimental settings are detailed in Appendix A.2.

Benchmarks. We evaluate our approach on three chart understanding benchmarks: ChartQA, ChartQAPro, and VisDoTQA. ChartQA and ChartQAPro span both human-authored and machine-generated questions, diverse chart types, and a wide range of reasoning tasks. VisDoTQA, constructed using our proposed framework, focuses on perceptual reasoning and serves as a challenging benchmark. Full dataset details are provided in Appendix A.3.

Evaluation metrics. We follow the official evaluation protocols for ChartQA and ChartQAPro, and use Relaxed Accuracy (RA) for VisDoTQA. Detailed evaluation metrics are provided in Appendix A.3.

4.2 Main Result

Table 3 summarizes our results on ChartQA, ChartQAPro, and VisDoTQA. Fine-tuning InternVL with VisDoTQA (VisDoT_InternVL) leads to consistent improvements over the InternVL baselines. On ChartQA, VisDoT_InternVL-2B and -4B achieve relative gains of +4.4% and +9.3%, reaching performance comparable to GPT-4o

Model	Param	ChartQA			ChartQAPro						VisDoTQA				
		Mac.	Hum.	Avg.	Factoid	Multi.	Conv.	Check.	Hypo.	Avg.	Pos.	Len.	Pat.	Ext.	Avg.
GPT-4o	–	–	–	85.7	35.76	46.72	34.75	45.49	28.91	37.67	74.57	55.00	36.33	57.03	57.14
Gemini-Flash-2.0	–	90.00	80.24	85.12	43.43	60.28	40.25	67.62	24.47	46.85	77.71	73.75	41.57	50.95	61.96
ChartGemma	3B	90.80	69.52	80.16	6.86	0.00	16.00	1.22	6.53	24.86	6.84	24.86	47.92	27.72	30.8
InternVL	2B	83.52	64.80	74.16	13.86	10.74	14.02	45.90	18.92	17.81	44.29	54.17	19.10	17.87	34.20
InternVL	4B	84.00	66.16	75.08	25.81	10.98	15.85	47.13	25.07	32.11	58.00	59.17	25.47	27.38	43.30
VisDoT_InternVL	2B	84.48 / 84.00 [†]	72.72	78.60 / 78.36 [†]	23.40	29.91	12.03	37.30	29.59	24.35	76.00	64.17	49.81	61.98	63.93
VisDoT_InternVL	4B	88.40 / 90.88[†]	80.32	84.36 / 85.60 [†]	32.75	47.20	15.38	53.28	40.82	34.54	87.14	69.58	67.42	77.95	76.52

Table 3: Comparison of baseline and VisDoT-tuned models across ChartQA, ChartQAPro, and VisDoTQA. VisDoT models leverage perception-following datasets and DoT-based supervision. For entries reported as *default / short-answer*[†], the value after “/” is obtained under the short-answer (final-answer-only) setting. Abbreviations: Mac. = Machine, Hum. = Human, Avg. = Average, Multi. = Multi-choice, Conv. = Conversational, Check. = Fact-checking, Hypo. = Hypothetical, Pos. = Position, Len. = Length, Pat. = Pattern, Ext. = Extract.

Method	# QA	ChartQA	ChartQAPro	VisDoTQA
ChartGemma	163K	83.96	36.61	26.79
ECD	320K	75.24	17.24	58.30
VisDoT (Ours)	7.4K	84.08	34.02	70.00

Table 4: Comparison of chart distillation/data-synthesis pipelines using publicly released instruction data: ChartGemma (Masry et al., 2025b) and ECD (Yang et al., 2025). All methods are fine-tuned on the same backbone (InternVL2.5-4B) and evaluated on ChartQA, ChartQAPro, and VisDoTQA. # QA denotes the number of QA pairs used for fine-tuning; for VisDoT, we randomly sample 7.4K QA pairs for this comparison, while Table 3 reports results using the full VisDoTQA training set (331,969 QA pairs).

and Gemini-Flash-2.0. The gains are more pronounced on the Human split (+7.9% and +14.2%). When constrained to short-answer output, VisDoT_InternVL-4B achieves 90.88% (+6.9%), establishing a new state of the art. Appendix C.1 further illustrates the importance of removing unnecessary reasoning steps in OCR-style questions. On ChartQAPro, VisDoT_InternVL-2B achieves 24.35% (+6.5%), while the 4B variant reaches 34.54% (+2.4%). Notable gains are observed in fact-checking (53.28%, +7.8%), multi-choice (47.20%), and hypothesis reasoning (+11.9% and +16.4%), surpassing GPT-4o. Despite limited training image diversity, the models generalize well to 157 unseen domains and unfamiliar chart layouts, demonstrating the effectiveness of perception-oriented tuning for domain generalization. As shown in Appendix C.2, baseline models often fail on charts with novel structures or multi-object configurations, whereas VisDoT_InternVL successfully handles these cases via DoT-based decomposition. Finally, on VisDoTQA, VisDoT_InternVL-4B improves the average accuracy from 43.30% to 76.52% (+33.2%), outper-

forming GPT-4o (+19.4%) and Gemini (+14.6%). The most significant gains appear in the Position (+29.1%), Length (+10.4%), Pattern (+42.0%) tasks, and Extract (+49.6%). These findings confirm that DoT-based perception–logic decomposition substantially strengthens visual grounding and compositional reasoning, enabling mid-sized LVLMs to match or exceed large closed models on both chart-specific and general visual question answering benchmarks. Additional model variants presented in Appendix B.1 exhibit consistent trends.

Comparison with chart distillation/data-synthesis pipelines. In addition to model baselines, we compare our approach with two representative chart distillation/data-synthesis pipelines that release instruction data publicly: ChartGemma (Masry et al., 2025b) and ECD (Yang et al., 2025). Table 4 summarizes the results under the same student backbone. Despite using substantially fewer QA pairs (7.4K vs. 163K/320K), VisDoT achieves the best performance on ChartQA (84.08%) and VisDoTQA (70.00%). Notably, VisDoTQA shows a large margin over ChartGemma (+43.21%) and ECD (+11.70%), suggesting that perception-following supervision paired with DoT-based decomposition is significantly more effective for perceptual reasoning than scaling synthetic instruction data alone. On ChartQAPro, VisDoT remains competitive (34.02%), outperforming ECD by a wide margin (+16.78%) while approaching ChartGemma (36.61%). These results indicate that VisDoT yields strong generalization across chart understanding benchmarks, particularly on perceptual-heavy tasks, while requiring orders of magnitude fewer fine-tuning QA pairs.

Method	ChartQA			VisDoTQA				
	Machine	Human	Avg.	Extract	Position	Length	Pattern	Avg.
Base	84.00	66.16	75.08	58.00	59.17	25.47	27.38	43.30
CoT-trained model	73.96 / 87.68 [†]	63.92	68.94 / 75.8 [†]	57.14	52.50	30.34	56.27	49.55
Perceptual-trained model	91.20	72.56	81.88	65.43	62.92	22.10	28.14	45.80
DoT-trained model	89.68 / 91.28 [†]	76.24	82.96 / 83.76 [†]	68.86	66.25	51.31	66.54	63.57
Perceptual + DoT (Ours)	83.44 / 91.20 [†]	76.96	80.20 / 84.08 [†]	78.57	65.83	61.42	71.10	70.00

Table 5: Performance comparison of training strategies on ChartQA and VisDoTQA using InternVL2.5-4B as the base model. [†] denotes short-answer evaluation (final-answer-only).

Model	Pos.	Len.	Pat.	Ext.	Avg.
All	83.43%	63.33%	63.30%	74.90%	72.32%
w/o Position	73.43%	66.25%	63.30%	72.24%	69.20%
w/o Length	85.14%	65.83%	64.04%	70.34%	72.50%
w/o Pattern	84.00%	64.58%	57.68%	69.96%	70.27%

Table 6: Performance across four perceptual tasks under ablation settings. For brevity, we denote Position as **Pos.**, Length as **Len.**, Pattern as **Pat.**, Extract as **Ext.**, and Average as **Avg.**.

4.3 Ablation Study: Effect of Perceptual Alignment and Thought Decomposition

Table 5 compares four training regimes on ChartQA and VisDoTQA, each using 7.4K matched samples (see Appendix A.4 for details). CoT reduces average accuracy on ChartQA, particularly on machine-generated questions, but regains a small margin when answers are forced into short form. On VisDoTQA, the same template raises accuracy, confirming its value for logic-only tasks, yet it also reveals sensitivity to question difficulty (Section 4.5). Perceptual yields the largest gains on ChartQA and moderate gains on VisDoTQA, indicating strong visual alignment but limited reasoning capacity. The DoT improves both datasets and remains stable across shifts in difficulty. On the ChartQA machine data, it avoids the -10.0% drop seen with CoT and instead scores +7.9% (free form) and +8.7% (short form). Perceptual + DoT maintains the Perceptual gains while adding DoT robustness, giving +9.0% on ChartQA and +20.3% on VisDoTQA, a 3.2 $\times$ larger boost than CoT on the latter. These results show that visual alignment and decomposition are complementary. Normalising answers to short form further sharpens performance, and the combined Perceptual + DoT strategy delivers reliable improvements across both OCR-style extraction and compositional reasoning tasks.

Dataset	Direct	CoT	Ours (DoT)
POPE	84.29%	84.64%	86.07%
MMMU	17.7%	35.5%	37.7%

Table 7: Zero-shot performance comparison on POPE and MMMU benchmarks using Direct prompting, CoT, and our proposed DoT strategy.

4.4 Ablation Study: Role of Individual Perceptual Task

Table 6 reports the results of fine-tuning VisDoT while sequentially removing each perceptual task from the training data. For fair comparison, 30K samples per task were randomly selected from VisDoTQA. Removing Position data produces the steepest decline: the average drops from 72.32% to 69.20% (- 3.12%), Position from 83.43% to 73.43% (- 10.00%), and Extract from 74.90% to 72.24% (- 2.66%). Position cues, therefore, support both spatial alignment and downstream extraction. Removing the Length data has a negligible or positive impact. w/o-Length scores 65.83% on Length tests (+2.50%), and w/o-Position scores 66.25% (+2.92%). On Position tests, w/o-Length attains 85.14% (+1.71%). Position thus substitutes for Length and reveals feature competition consistent with cognitive findings. More data alleviate this competition (Table 3).

Overall, Position is critical for stable accuracy and extraction, Length is partly redundant, and Pattern governs categorical discrimination. Perception-following fine-tuning with ample samples remains essential for reliable reasoning on visualized data.

4.5 Cross-Domain Evaluation of DoT Prompt’s Robustness

Table 7 evaluates DoT under zero-shot transfer using InternVL2.5-78B on POPE and MMMU. POPE targets natural object images, whereas MMMU spans six academic domains and 30 visual for-

mats. DoT exceeds both Direct and CoT baselines: +1.43% on POPE and +2.2% over CoT on MMMU. These findings suggest that DoT is not a simple variant of CoT but a structured reasoning strategy grounded in the separation of perception and logic. The consistent improvements across natural image QA and diverse visual domains indicate that DoT has strong potential for domain-invariant transfer. Furthermore, DoT’s explicit separation of perceptual and logical sub-questions enables stable reasoning over decomposed sub-queries, supporting its role as a general-purpose reasoning strategy for LVLMs. Moreover, DoT satisfies both components of the structured reasoning formulation in Equation 2, enabling more robust and generalizable visual reasoning beyond chart-centric settings. Detailed experimental results are provided in Appendix B.2.

5 Conclusion

This work introduced VisDoT, a framework for visual question answering on data visualizations. Grounded in graphic perception theory, we formalized four perception-following task families (Position, Length, Pattern, and Extract) and designed a DoT prompting strategy that first performs perceptual grounding and then logical reasoning. VisDoT enables robust visual grounding even with compact LVLMs and produces interpretable step-wise reasoning traces. We also release a perception-following dataset that embodies these design principles. Experiments across diverse visualization types show that current LVLMs still face substantial challenges in understanding visualizations and that grounding accuracy is a primary bottleneck in visual reasoning. VisDoT mitigates this limitation by explicitly integrating grounding with reasoning through the DoT prompt. Future work will extend this approach to other grounding-intensive domains and will leverage DoT reasoning traces to analyze whether LVLm failures stem from deficits in perception or reasoning.

Limitations

While the current study primarily evaluated the framework using raster chart formats (e.g., PNG, JPEG), future work could further investigate its applicability to vector-based formats (e.g., SVG, PDF) that preserve object boundaries, textual metadata, and hierarchical group information, poten-

tially enabling richer grounding capabilities. Extending the VisDoT framework to more complex visualization layouts, such as multi-panel dashboards with interactive elements, may provide additional insights. Moreover, although VisDoT focuses on perception tuning for enhanced visual grounding, incorporating instruction tuning could better address real-world queries involving higher-order discourse structures (e.g., comparison, explanation, prediction). Future research could explore hybrid training strategies that jointly optimize visual token recognition and discourse-aware response quality, possibly through curriculum-based approaches interleaving perception-first and instruction-rich samples. Finally, prompt design for DoT remains an open area for refinement, and further investigations are needed to identify strategies that can maximize the effectiveness and generalizability of the framework.

Acknowledgments

This research was supported by the MSIT(Ministry of Science and ICT), Korea, under the ITRC(Information Technology Research Center) support program(IITP-2026-RS-2020-II201789), and the Artificial Intelligence Convergence Innovation Human Resources Development(IITP-2026-RS-2023-00254592) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation).

References

- Saeed Amizadeh, Hamid Palangi, Alex Polozov, Yichen Huang, and Kazuhito Koishida. 2020. Neuro-symbolic visual reasoning: Disentangling. In *International Conference on Machine Learning*, pages 279–290. Pmlr.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. 2019. Scene text visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4291–4301.
- Victor Cărbune, Hassan Mansoor, Fangyu Liu, Rahul Aralikkatte, Gilles Baechler, Jindong Chen, and Abhanshu Sharma. 2024. Chart-based reasoning: Transferring capabilities from llms to vlms. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 989–1004.

- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, and 1 others. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- William S Cleveland and Robert McGill. 1984. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association*, 79(387):531–554.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Yucheng Han, Chi Zhang, Xin Chen, Xu Yang, Zhibin Wang, Gang Yu, Bin Fu, and Hanwang Zhang. 2023. [Chartllama: A multimodal llm for chart understanding and generation](#). *Preprint*, arXiv:2311.16483.
- Ziniu Hu, Ahmet Iscen, Chen Sun, Kai-Wei Chang, Yizhou Sun, David Ross, Cordelia Schmid, and Alireza Fathi. 2023. Avis: Autonomous visual information seeking with large language model agent. *Advances in Neural Information Processing Systems*, 36:867–878.
- Zaid Khan, Vijay Kumar BG, Samuel Schuster, Manmohan Chandraker, and Yun Fu. 2023. Exploring question decomposition for zero-shot vqa. *Advances in Neural Information Processing Systems*, 36:56615–56627.
- Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvasi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. 2023. Pix2struct: Screen-shot parsing as pretraining for visual language understanding. In *International Conference on Machine Learning*, pages 18893–18912. PMLR.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Evaluating object hallucination in large vision-language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305.
- Jingming Liu, Yumeng Li, Boyuan Xiao, Yichang Jian, Ziang Qin, Tianjia Shao, Yao-Xiang Ding, and Kun Zhou. 2025. [Autonomous imagination: Closed-loop decomposition of visual-to-textual conversion in visual reasoning for multimodal large language models](#). *Preprint*, arXiv:2411.18142.
- Peiyu Liu, Ze-Feng Gao, Xiao Zhang, Wayne Xin Zhao, and Ji-Rong Wen. 2024. [Enhancing parameter-efficient fine-tuning with simple calibration based on stable rank](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6024–6035, Torino, Italia. ELRA and ICCL.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279.
- Ahmed Masry, Mohammed Saidul Islam, Mahir Ahmed, Aayush Bajaj, Firoz Kabir, Aaryaman Kartha, Md Tahmid Rahman Laskar, Mizanur Rahman, Shadikur Rahman, Mehrad Shahmohammadi, Megh Thakkar, Md Rizwan Parvez, Enamul Hoque, and Shafiq Joty. 2025a. [ChartQAPro: A more diverse and challenging benchmark for chart question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19123–19151, Vienna, Austria. Association for Computational Linguistics.
- Ahmed Masry, Mehrad Shahmohammadi, Md Rizwan Parvez, Enamul Hoque, and Shafiq Joty. 2024. Chartinstruct: Instruction tuning for chart comprehension and reasoning. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 10387–10409.
- Ahmed Masry, Megh Thakkar, Aayush Bajaj, Aaryaman Kartha, Enamul Hoque, and Shafiq Joty. 2025b. [ChartGemma: Visual instruction-tuning for chart reasoning in the wild](#). In *Proceedings of the 31st International Conference on Computational Linguistics: Industry Track*, pages 625–643, Abu Dhabi, UAE. Association for Computational Linguistics.
- Fanqing Meng, Wenqi Shao, Quanfeng Lu, Peng Gao, Kaipeng Zhang, Yu Qiao, and Ping Luo. 2024. [Chartassistant: A universal chart multimodal language model via chart-to-table pre-training and multitask instruction tuning](#). *Preprint*, arXiv:2401.02384.
- OECD. 2024. Organisation for economic co-operation and development (oecd). <https://www.oecd.org>. Accessed: 2025-07-28.
- OpenAI. 2024. Gpt-4o system card. <https://openai.com/index/gpt-4o-system-card/>. Accessed: 2025-07-28.
- Our World in Data. 2024. Our world in data: Research and data to make progress against the world’s largest problems. <https://ourworldindata.org>. Accessed: 2025-07-28.
- Pew Research Center. 2024. Pew research center | nonpartisan, non-advocacy public opinion polling. <https://www.pewresearch.org>. Accessed: 2025-07-28.

- Statista. 2024. Statista: Statistics and studies from more than 22,500 sources. <https://www.statista.com>. Accessed: 2025-07-28.
- Stanley Smith Stevens. 2017. *Psychophysics: Introduction to its perceptual, neural and social prospects*. Routledge.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivi re, and 1 others. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Alexander Vogel, Omar Moured, Yufan Chen, Jiaming Zhang, and Rainer Stiefelhagen. 2025. Refchartqa: Grounding visual answer on chart images through instruction tuning. *arXiv preprint arXiv:2503.23131*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Ziyue Wang, Chi Chen, Peng Li, and Yang Liu. 2023. Filling the image information gap for vqa: Prompting large language models to proactively ask questions. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2874–2890.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Qiong Wu, Xiangcong Yang, Yiyi Zhou, Chenxin Fang, Baiyang Song, Xiaoshuai Sun, and Rongrong Ji. 2025. Grounded chain-of-thought for multimodal large language models. *arXiv preprint arXiv:2503.12799*.
- Yuwei Yang, Zeyu Zhang, Yunzhong Hou, Zhuowan Li, Gaowen Liu, Ali Payani, Yuan-Sen Ting, and Liang Zheng. 2025. Effective training data synthesis for improving mllm chart understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.
- Haowei Zhang, Jianzhe Liu, Zhen Han, Shuo Chen, Bailan He, Volker Tresp, Zhiqiang Xu, and Jindong Gu. 2024a. Visual question decomposition on multimodal large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1926–1949.
- Yizhe Zhang, He Bai, Ruixiang Zhang, Jiatao Gu, Shuangfei Zhai, Josh Susskind, and Navdeep Jaitly. 2024b. How far are we from intelligent visual deductive reasoning? *arXiv preprint arXiv:2403.04732*.

A Detailed Descriptions

A.1 Perceptual Task Details

As briefly introduced in Section 3.1, we define four perceptual tasks for visualized data interpretation: Position, Length, Pattern and Extract. These tasks are based on foundational principles of graphical perception (Cleveland and McGill, 1984). This section provides a more detailed explanation of the rationale behind selecting these perceptual tasks.

Based on the prior work of Cleveland and McGill (Cleveland and McGill, 1984), we select three core perceptual tasks—**Position**, **Length**, and **Pattern**—as foundational operations for visual reasoning. Their study identified ten elementary perceptual tasks essential for chart interpretation and ranked them according to their perceptual accuracy. Among these, Position demonstrated the highest interpretive accuracy in human studies and is therefore included as a primary visual task. Length belongs to the second-highest accuracy tier and, according to Stevens’ power law (Stevens, 2017), is one of the least distorted visual encodings. In contrast, other candidates, such as angle, area, and volume, were excluded due to their known susceptibility to perceptual bias and overestimation. The task pattern—while not suitable for precise quantitative estimation—is included for its effectiveness in categorical differentiation, as emphasized by (Cleveland and McGill, 1984). The **Extract** task refers to directly retrieving explicitly labeled numerical values from the visual input. For instance, to answer the question “What is the difference in GDP between the United States and Canada?”, the model must directly extract the GDP values of both countries from the chart.

In summary, we define the following four perceptual tasks as the fundamental components for visual reasoning over visualized data: Position, Length, Pattern, and Extract.

A.2 Training Configurations

We fine-tuned four LVLMS under different training strategies, including parameter-efficient LoRA-based tuning and full fine-tuning. The specific configuration for each model is detailed below, and its corresponding hyper-parameters are summarized in Table 8. InternVL2.5-2B and InternVL2.5-4B

Model	Epoch	Learning Rate	Batch Size	Hours
InternVL-2B	1	4e-5	16	10
InternVL-4B	1	4e-5	10	12
Gemma3-4B	1	2e-4	4	20
Qwen2.5-VL-3B	1	2e-4	64	29

Table 8: Training configurations used for different LVLM backbones, including number of epochs, learning rate, batch size, and training time.

were fine-tuned using the LoRA method (rank = 16), applied exclusively to the language decoder. Both the vision encoder and the multi-modal projector were kept frozen during training. Gemma3-4B was fine-tuned using the Unsloth framework in a parameter-efficient manner. The vision encoder was frozen, and LoRA adapters (rank = 8, $\alpha = 16$) were applied to the language model. Qwen2.5-VL-3B was fine-tuned using a full fine-tuning strategy with the exception of the vision encoder, which remained frozen. Both the multi-modal MLP and the language model components were updated during training to adapt the model to downstream tasks. We used six A6000 GPUs, and the total training took approximately 71 hours.

A.3 Benchmarks, Metrics and Models Details

Benchmarks We evaluate our approach on three chart understanding benchmarks. ChartQA consists of two subsets: ChartQA-Human, which contains human-authored questions, and ChartQA-Machine, which includes questions generated from chart summaries. ChartQAPro is built from unseen charts collected across 157 domains, covering not only simple bar and line charts but also complex visual structures such as multi-chart layouts, stacked bar charts, dashboards, and infographics. It includes five reasoning types: factoid, conversational, multiple-choice, hypothetical, and fact-checking. We further evaluate on VisDoTQA, which is constructed using the VisDoT framework with images that are not used in training. VisDoTQA comprises questions incorporating four types of perceptual information and requires multi-object perception, grounding, and advanced reasoning, providing a challenging benchmark for comprehensive chart understanding.

Evaluation Metrics For ChartQA, we adopt the Relaxed Accuracy (RA) metric for evaluation. For ChartQAPro, evaluation is performed as follows. Numerical answers are considered correct if they

fall within a $\pm 5\%$ error margin. For year values, an exact match (EM) is required. Textual answers are evaluated using the average normalized levenshtein similarity (ANLS (Biten et al., 2019)) score. Multiple-choice (MCQ) and fact-checking questions are assessed using the EM criterion. For VisDoTQA, we employ the RA metric for evaluating answers.

Models We fine-tuned Gemma3-4B (Apache 2.0 license) using the Unsloth framework (Apache 2.0 license), applying LoRA (rank = 8, $\alpha = 16$) to the text decoder. Qwen2.5-VL-3B (Apache 2.0 license) was fully fine-tuned except for the frozen vision encoder, and InternVL2.5-2B/4B (MIT license) models were fine-tuned using LoRA (rank = 16) applied only to the language decoder. For baseline inference, we used the Ollama platform (MIT license). All models were trained and evaluated on the ChartQA and ChartQAPRO datasets (MIT license). We also referred to the ChartGemma implementation (GPL-3.0 license) during preliminary experiments.

A.4 Detailed Description of the Ablation Study

To evaluate the effectiveness of our proposed Perceptual tasks and decomposition-of-thought (DoT) strategy in visualized data reasoning, we design five training configurations using the same image-question dataset. For each of the 7,400 image-question pairs, responses are generated under the following distinct paradigms:

- **Base:** Direct answers are generated without any intermediate reasoning structure or prompt engineering, serving as the most naive baseline.
- **CoT:** Responses are generated by reasoning through a step-by-step process before producing the final answer.
- **Perceptual:** Each sample includes explicit visual cues—such as spatial location, color, and order—prior to answering. While the output format remains consistent with Base, the model is guided to develop a perception-aligned reasoning mechanism.
- **DoT:** Unlike CoT, this configuration decomposes complex questions into a series of sub-questions. The model sequentially answers each sub-question to arrive at the final answer.

- **Perceptual + DoT (Ours):** This is the most structured setting, combining both perceptual grounding and compositional reasoning. Each instance involves a two-stage process: the Question Phase decomposes the original query into perceptual and cognitive sub-questions, and the Answer Phase sequentially generates intermediate answers leading to the final response.

These five configurations form the basis of the comparative experiments shown in Table 5, providing a comprehensive analysis of the impact of our Vis-DoT framework on the performance of LVLMs.

B Additional Results and Analysis

B.1 Additional Results using Other Models

In addition to the experimental results presented in Section 4.2, we provide supplementary results using alternative model backbones. Our results are shown in Table 9. These additional experiments consistently corroborate the findings reported in the main text, further validating the robustness and generalizability of our conclusions.

B.2 Detailed Results for POPE, MMMU

To validate the generalizability of our proposed DoT prompting strategy, we evaluate it on two external benchmarks in a zero-shot setting: POPE (Li et al., 2023), which targets hallucination detection in natural images, and MMMU (Yue et al., 2024), which spans diverse academic visual reasoning tasks.

We independently sample **280 questions each** from POPE and MMMU to ensure consistent and controlled comparisons across all prompting strategies.

Quantitative Results. As summarized in Table 6, DoT consistently outperforms both Direct and CoT prompting baselines across both benchmarks:

- **POPE:** DoT achieves **86.07%** accuracy, outperforming CoT (84.64%) and Direct (84.29%) by **+1.43%p** and **+1.78%p**, respectively.
- **MMMU:** DoT reaches **37.7%**, improving over CoT (35.5%) by **+2.2%p** and over Direct (17.7%) by a substantial **+20.0%p**.

These results demonstrate that DoT is more robust across both binary judgment tasks (POPE) and multi-format reasoning benchmarks (MMMU).

Why is DoT More Effective? Unlike CoT, which decomposes every question into multi-step reasoning regardless of complexity, DoT dynamically determines when reasoning is necessary based on perceptual cues. This perception-first principle allows DoT to avoid over-decomposition and reduce hallucinations, especially when visual evidence alone suffices.

- In **POPE**, CoT sometimes hallucinates object presence due to speculative reasoning (e.g., “There might be a stool beside the plant”). DoT, in contrast, first verifies perceptual absence (e.g., “No object resembling a chair is present”) before concluding.
- In **MMMU**, CoT often misinterprets axis semantics or misaligns chart elements due to rigid, uniform reasoning chains. DoT first inspects structural layout (e.g., axis labels, bar heights), and then performs reasoning grounded in perceptual understanding.

POPE Example:

- **Question:** Is there a chair in the image?
- **Answer:** No
- **Direct:** No chair visible. (Correct)
- **CoT:** There is a stool beside the plant. (Incorrect – hallucination)
- **DoT:** No object resembling a chair (with back and legs) is present. (Correct)

MMMU Example:

- **Question:** What circumstances lead to the typical growth form of the trees shown in the image?
- **Answer:** D
- **Direct:** r (Incorrect – uninformative)
- **CoT:** Focuses on bark pattern but overlooks canopy structure. (Incorrect)
- **DoT:** Inspects overall morphology and spacing, then selects D. (Correct)

These examples demonstrate that DoT not only improves accuracy but also enhances interpretability and factual consistency over existing prompting strategies.

Model	Param	ChartQA			ChartQAPro						VisDoTQA				
		Mac.	Hum.	Avg.	Factoid	Multi.	Conv.	Check.	Hypo.	Avg.	Pos.	Len.	Pat.	Ext.	Avg.
Gemma (CoT)	4B	–	58.24	–	11.01	24.77	–	47.13	12.24	–	46.86	38.75	17.98	43.35	37.41
Qwen (CoT)	3B	–	67.12	–	24.46	32.24	5.40	45.49	33.97	–	67.43	54.58	48.69	65.78	59.82
VisDoT_Gemma	4B	84.16	77.04	80.60	23.40	28.04	–	54.10	38.78	–	79.14	58.75	59.18	70.72	67.32
VisDoT_Qwen	3B	86.96	79.60	83.28	26.46	37.38	–	43.44	40.82	–	77.14	55.42	59.55	73.00	68.04

Table 9: Comparison of CoT and VisDoT-tuned models across ChartQA, ChartQAPro, and VisDoTQA. VisDoT models leverage perception-following datasets and DoT-based supervision. Gray cells (if any) denote short-answer performance. Abbreviations: Mac. = Machine, Hum. = Human, Avg. = Average, Multi. = Multi-choice, Conv. = Conversational, Check. = Fact-checking, Hypo. = Hypothetical, Pos. = Position, Len. = Length, Pat. = Pattern, Ext. = Extract.

C Case Studies

C.1 How Unnecessary Reasoning Steps Degrade Performance?

In OCR-style tasks involving relatively simple questions, directly providing short answers proves to be more effective than generating extended reasoning traces. As shown in Figures 5 and 6, unnecessary decomposition or step-by-step reasoning may introduce errors or distract the model from the correct answer, ultimately degrading performance. These findings highlight that, for low-complexity visual queries, minimal reasoning with concise output formatting can lead to superior accuracy.

C.2 Advantages of DoT Prompting

Base models often fail to reason over visualized data involving unseen visual structures or multi-object references. In contrast, models trained with VisDoTQA effectively utilize the DoT-style question decomposition strategy to derive the correct answers. A representative example is shown in Figure 7.

C.3 Evaluation Case Studies

In this section, we present case studies demonstrating how fine-tuned models effectively respond to chart-based questions. Representative examples are selected from the ChartQA, ChartQAPro, and VisDoTQA human split subsets. These examples highlight the enhanced reasoning capabilities achieved through fine-tuning.

Our analysis focuses on the outputs of fine-tuned models, as baseline models often fail to produce verifiable or coherent answers in complex chart reasoning tasks. Each example illustrates the model’s ability to accurately recognize visual cues and perform multi-step reasoning to arrive at the correct answer.

Figure 8, 9: Results on the ChartQA Human task

Figure 10, 11, 12, 13: Results on the ChartQAPro task

Figure 14, 15, 16, 17: Results on the VisDoTQA task

To better understand how models handle complex chart-based questions, we analyzed responses on VisDoTQA using a unified prompt template. This comparison reveals clear behavioral differences between baseline and fine-tuned models.

The InternVL2.5 baseline often generated short, direct answers and failed to follow the multi-step reasoning structure provided in the prompt. Even when structured guidance was available, the model tended to return a single numerical value without explaining the reasoning process. This indicates difficulties in comprehending complex questions and producing accurate answers.

The Qwen2.5-VL-3B baseline showed a stronger tendency toward step-by-step reasoning, aligning with its general instruction-following capabilities. However, it frequently misinterpreted visual elements (e.g., bar heights or values) and exhibited weak numerical reasoning, leading to unit confusion or incorrect calculations. These errors suggest that while the model attempts structured inference, it lacks the visual grounding and numerical precision required for chart-based reasoning.

In contrast, all models fine-tuned on the VisDoTQA dataset consistently followed the reasoning structure outlined in the prompt. Despite using the same system prompt as the baselines, the fine-tuned models successfully decomposed complex questions, accurately interpreted visual content, and generated correct answers through sequential reasoning. This demonstrates that VisDoT-based fine-tuning significantly improves a model’s ability to align visual perception with logical inference, even without additional prompt engineering.

D Detailed Prompt

D.1 Perception-following Question Generation

Below are the prompt templates used for generating Perception-following Questions. Each template is designed to elicit questions corresponding to one of the four perceptual tasks: Position, Length, Pattern, and Extract. The full prompt is presented in Table 10, 11, 12 and 13.

D.2 Decomposition-of-Thought Prompt

We present the full DoT prompting template referenced in Section 3.2. The prompt first instructs the model to decompose a given question into sub-questions, explicitly prioritizing the generation of perception-oriented sub-questions before cognitive ones. Subsequently, the model is guided to sequentially generate intermediate answers for each sub-question, culminating in a final answer that integrates the accumulated reasoning steps. The full prompt is presented in Table 14.

D.3 Model Inference Prompt

We present the inference prompts used across various benchmarks and experimental settings. All prompts were provided to the model as system-level instructions. The full prompt is presented in Table 15.

E Representative Examples of the VisDoTQA Dataset

To qualitatively illustrate the diversity and reasoning patterns captured by our VisDoTQA dataset, we present four representative examples, each corresponding to a distinct type of reasoning: **Position-based**, **Length-based**, **Pattern-based**, and **Extractive**.

Data Source. All chart images used in this appendix are sourced from the publicly available **ChartQA** dataset (Masry et al., 2022). The visual content remains unchanged, while the questions and answers are newly constructed based on our proposed DoT-guided reasoning framework.

(1) Position-based Reasoning

- **Question:** Rank the regions based on their dark purple bar revenue, and determine how much higher the top-ranked region’s revenue is compared to the bottom-ranked region.

- **Answer:** The dark purple (2020) revenues are: The Americas (565,023), Europe (326,613), Japan (301,187), Other (115,694). The difference between the top and bottom regions is $565,023 - 115,694 = 449,329$.

Figure 18 illustrates this position-based reasoning example.

(2) Length-based Reasoning

- **Question:** Find the median time for trajectory in black between 2010 and 2019, and compare it to the median time for blue trend.
- **Answer:** For the black line (Women), the median is 233.5 minutes; for the blue line (Men), it is 203 minutes. Thus, the black trend is longer by **30.5 minutes**.

Figure 19 shows the corresponding length-based reasoning instance.

(3) Pattern-based Reasoning

- **Question:** Identify the group that holds about one-fourth of the total seats occupied by the five largest groups combined.
- **Answer:** The total seats of top-5 groups = 566. One-fourth is approximately 141.5. The Center-left (S&D) group holds 146 seats, which is closest. **Answer: Center-left (S&D).**

Figure 20 demonstrates the pattern-based reasoning task.

(4) Extractive Reasoning

- **Question:** Calculate the difference between the highest and lowest percentage categories.
- **Answer:** Highest = 53% (Regular user), Lowest = 4% (Seasonal user). Difference = $53 - 4 = 49\%$.

Figure 21 presents the extractive reasoning case.

System Prompt: Question Generation - Position Prompt

You are an AI assistant that generates **location-based analytical questions** from structured visual data, such as charts and tables. Your task is to analyze the image and generate **questions that extract values based on their spatial position**.

How to Generate Questions

- Use **the relative position of elements** instead of direct category names as the primary reference.
- However, **also include the actual name/label in parentheses** after each positional reference for clarity.
- Example:
 - **Rows (Vertical Order)**: "third region from the top (South America)", "second from the bottom (Europe)"
 - **Columns (Horizontal Order)**: "leftmost column (2019)", "second from the right (2021)"
- Ensure a mix of different **positional references** (top, bottom, upper, lower, first, last, left, right).
- Avoid simple retrieval questions—each question should require some form of reasoning (comparison, sum, difference, etc.).

Task

Generate **2 complex analytical questions** based on spatial positioning.
Use both **relative positions** and **actual names/labels** in parentheses.

Table 10: Perception-following Question Generation - Position Prompt

System Prompt: Perception-following Question Generation - Length Prompt

System Prompt

You are an expert in analyzing data visualizations and generating insightful questions.
Based on a given chart image, generate appropriate questions.

Required Operations


- **extract**: Ask for the value of a specific category.
- **ranking**: Identify the item, year, or group with the highest, 2nd, 3rd,... lowest value.
- **comparison**: Compare two or more values or categories.
- **counting**: Count how many items meet a specific condition.
- **addition & subtraction**: Summing or finding differences between multiple values.
- **multiplication & division**: Calculating proportional changes or percentage comparisons.
- **ratio**: Determining how many times one value is larger or smaller than another.
- **rate of change**: Identify intervals or categories based on how much their values have changed over time.

Output Format

"type": "extract", "question": "What was the export share in 2007?",
"type": "ranking", "question": "Which month had the lowest hiring rate?",
"type": "comparison", "question": "Compare the hiring rates of May 2020 and February 2021.",
"type": "comparison", "question": "Which country data is consistently above 1.5 kg?",
"type": "rate of change", "question": "Which two consecutive years showed the largest change?",
"type": "subtraction", "question": "How did the export share change from 2006 to 2009?",
"type": "comparison + counting", "question": "How many years had a higher export share than 2010?",
"type": "subtraction + ranking", "question": "What is the difference between the highest and lowest export share years?"

Table 11: Perception-following Question Generation - Length Prompt

Image:



Query:

What was the annual revenue from beer excise duty in Denmark in 2019?

Short Answer:

"110.06" (correct)

CoT Answer:


To determine the annual revenue from beer excise duty in Denmark for 2019, we need to analyze the provided bar graph. The graph displays the excise duty revenue in million euros from 2008 to 2016. \n\n**Key Observations:** \n\n

The x-axis represents the years, while the y-axis shows the revenue. \n\n Each year has a corresponding blue bar indicating the revenue for that year. \n\n Notably, there is no data point for 2017 or 2018 on the graph. \n\n**Deriving the Answer:** \n\n Given the absence of data for 2017 and 2018, we cannot directly calculate the annual revenue from beer excise duty in Denmark for these years using the provided graph. However, if we were to extrapolate based on the trend shown (assuming it continues), we would look at the last available data point, which is 2016. \n\n

\n\n**Conclusion:** \n\n Without explicit data for 2017 and 2018, we cannot accurately determine the annual revenue from beer excise duty in Denmark for these years using the provided graph alone. The graph only offers information up to 2016, and without additional context or data points beyond that year, any conclusion about 2017 and 2018 would be speculative. \n\n### Answer: Unknown

Figure 5: Short answer case 1

Image:



Query:

What is one of the staple foods of the south Asian country?

Short Answer:

"Rice" (correct)

CoT Answer:

Let's break down this problem.

\n\n**Question Phase:**\n1. What is one of the staple foods of the south Asian country?\n2. Which crop has a production volume of 110 million metric tons?\n\n**Solution Phase:**\n1) The crop with a production volume of 110 million metric tons is Wheat.\n2) Therefore, one of the staple foods of the south Asian country is Wheat.

\n\n**Answer:** Wheat

Figure 6: Short answer case 2

System Prompt: Perception-following Question Generation - Pattern Prompt

You are an AI assistant that generates analytical questions based on recurring visual patterns in charts. Your task is to create **complex, meaningful questions** that require pattern recognition and reasoning over visual groupings such as colors, legends, markers, or labeled categories, rather than simple label-based lookup.

Required Operations (use each at least twice across 20 questions):

- **Pattern Counting:** Identifying how many elements share a common visual pattern (e.g., "How many wedges are in shades of blue?").
- **Color/Legend Matching:** Comparing values between categories based on shared legend color or visual symbol.
- **Pattern-based Addition & Subtraction:** Summing or comparing values across elements sharing the same visual style (e.g., line types, bar textures).
- **Ratio:** Determining the proportional relationship between visually grouped categories.
- **Comparison:** Identifying which visually marked group has greater or lesser values.
- **Pattern Ranking:** Ordering patterns by magnitude or frequency of values (e.g., which color category has the 2nd highest value).

Important Guidelines

- Clearly include **at least two reasoning operations** per question.
- **All listed operations** must appear at least **twice across the 20 questions**.
- Avoid surface-level or pattern-mention-only questions—require true reasoning over the visual encodings.
- Focus on pattern-informed reasoning such as matching legends, interpreting repeated visual cues, or comparing across symbol groups.

Example Questions

1. What is the total percentage of the three categories shaded in red? (Pattern Counting, Addition)
2. Which legend color represents the category with the largest difference from the average? (Legend Matching, Comparison)
3. How much higher is the value for the striped bar than the dotted bar? (Pattern Comparison, Subtraction)
4. Which marker shape corresponds to the second-highest value? (Visual Pattern Ranking)
5. Among all categories with similar shades of green, which has the lowest count? (Color Grouping, Ranking)
6. How many segments share both the same color and bar direction? (Pattern Matching, Counting)

Table 12: Perception-following Question Generation - Pattern Prompt

Grenfell fire: how quickly it spread (east/north face of tower)



Statement : The fire had spread to more than twice as many flats by 2:53 am compared to 2:09 am.

InternVL2.5-2B (Base) : **True (Wrong)**

InternVL2.5-2B (**VisDoT**) : Let's break down this problem.

****Question Phase:****


1. What is the number of flats at 2:09 am?
2. What is the number of flats at 2:53 am?
3. Is the number of flats at 2:53 am more than twice the number of flats at 2:09 am?

****Solution Phase:****

- 1) According to the chart, there were ****34 flats**** at 2:09 am.
- 2) At 2:53 am, the number of flats was ****61****.
- 3) Twice the number at 2:09 am is **** $2 \times 34 = 68$ ****. Since ****61**** is less than ****68****, the number at 2:53 am is not more than twice the number at 2:09 am.

****Final Answer:**** False **(Correct)**

Figure 7: DoT adv example



Question : What is the difference between the space available in Europe?

Qwen2.5-VL-3B (Base) : The difference between the space available in Europe in 2011 and 2006 is **0.8 million square meters**.

Qwen2.5-VL-3B (VisDoT) : Let's break down this problem.

****Question Phase:****


1. What is the exhibition space in Europe for 2006?
2. What is the exhibition space in Europe for 2011?
3. What is the difference between the two values?

****Solution Phase:****

- 1) The exhibition space in Europe for ****2006**** is ****14.6 million square meters****.
- 2) The exhibition space in Europe for ****2011**** is ****15.6 million square meters****.
- 3) To find the difference, we subtract the value of 2006 from the value of 2011: ****15.6 - 14.6 = 1.0 million square meters****.

****Thus, the difference between the space available in Europe is 1.0 million square meters.****

Figure 8: ChartQA Human task 1



Question : What's the sum of the two middle bars in this chart?

Gemma3-4B (Base) : Okay, let's analyze the chart.

The two middle bars are:

- Brazil: 2918
- Pakistan: 1350

Adding those together, we get: $2918 + 1350 = 4268$

So, the sum of the two middle bars is **4268**.

Gemma3-4B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. What are the values of the two middle bars (Pakistan and Uzbekistan)?
2. What is the sum of these two values?

****Solution Phase:****

1) The chart shows the following values for the two middle bars:

- Pakistan: ****1350 thousand metric tons****
- Uzbekistan: ****762 thousand metric tons****

2) Calculating the sum:

- Sum = $1350 + 762 = \mathbf{**2112**}$

****Final Answer:**** The sum of the two middle bars is ****2112 thousand metric tons****.

Figure 9: ChartQA Human task 2

The £20.0 billion borrowed in September 2022 was £2.2 billion more than in September 2021 and £11.0 billion more than in pre-coronavirus (COVID-19) September 2019

Public sector net borrowing excluding public sector banks, £ billion, UK, September 2019 to September 2022



Publication: Public sector finances, UK: September 2022

Office for National Statistics

Statement : The public sector net borrowing excluding banks (psnb ex) for July 2020 was slightly over £20 billion.

Qwen2.5-VL-3B (Base) : The image shows data for January, May, September 2020, January 2021, May 2021, and September 2021. **It doesn't provide data** for July 2020. Therefore, it's impossible to determine if the public sector net borrowing excluding banks (psnb ex) for July 2020 was slightly over £20 billion. The answer is **unanswerable**.

Qwen2.5-VL-3B (VisDoT) : Let's break down this problem.

****Question Phase:****


1. What is the public sector net borrowing excluding banks (psnb ex) for July 2020?
2. Is it slightly over £20 billion?

****Solution Phase:****

- 1) Looking at the chart, the bar for July 2020 is slightly above the £20 billion mark on the Y-axis.
- 2) Therefore, the answer to the question is ****True****.

****Answer:** True**

Figure 10: ChartQAPro Fact Checking task



Question : what is the median of the percentages appeared on the bars in the chart?

Gemma3-4B (Base) : The chart shows percentages for several universities. The percentages are: 5%, 10%, 20%, 30%, 45%, and 50%. First, we sort the percentages in ascending order: 5%, 10%, 20%, 30%, 45%, 50%. The median is the middle value in the sorted list. Since there are 6 values, the median is the 3rd value, which is 20%. The answer is **20%**.

Gemma3-4B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. What are the percentages represented by each bar in the chart?
2. How do we calculate the median of these percentages?

****Solution Phase:****

1) The chart shows the following values:

- China (red): ****45%****
- U.S.A. (green): ****20%****
- South Korea (yellow): ****5%****
- Other (blue): ****30%****


2) To calculate the median, we first sort the values in ascending order: 5%, 20%, 30%, 45%.

Since there are four values (an even number), the median is the average of the two middle values:

$$(20 + 30) / 2 = \mathbf{**25%**}$$

****Final Answer:**** The median percentage is ****25%****.

Figure 11: ChartQAPro Factoid task



Question : what is the median of the percentages appeared on the bars in the chart?

Qwen2.5-3B (Base) : $181.3 - 175 = 6.3$

Qwen2.5-3B (VisDoT) : Let's break down this problem.

****Question Phase:****


1. What is the value of the Washington Area Home Price Index in April?
2. If the index were 10 points higher, what would the new value be?
3. How much greater would this new value be compared to the Composite-20 in April?

****Solution Phase:****

- 1) The Washington Area Home Price Index in April is ****181.3****.
- 2) If it were 10 points higher, the new value would be ****191.3**** ($181.3 + 10$).
- 3) The Composite-20 value in April is ****135.8****.
To compare, we subtract: ****191.3 - 135.8 = 55.5****

****Final Answer:** 55.5**

Figure 12: ChartQAPro Hypothetical task



Question : if the average play time for females who play in depth were 5 minutes lower, would women who play in depth play more than males in depth in age bracket 55-64?

InternVL2.5-2B (Base) : d) renewable

InternVL2.5-2B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. What is the energy efficiency of each energy source?
2. Which energy source has the highest energy efficiency?

****Solution Phase:****

- 1) According to the chart, the energy efficiencies are as follows:
 - Coal: ****44.58%****
 - Gas: ****45.02%****
 - Nuclear: ****45.77%****
 - Renewable: ****45.48%****

- 2) Among these, ****Nuclear**** has the highest energy efficiency at ****45.77%****.

****Final Answer:** c) Nuclear**

Gemma3-4B (Base) : The chart shows the energy efficiency of coal, gas, nuclear, and renewable sources. **Coal has 45.8%, gas has 2.0, nuclear has 2.2 and renewable has 2.5.** Therefore, renewable has the highest energy efficiency. The answer is **d**

Gemma3-4B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. Which energy source has the highest energy efficiency according to the chart?
2. What is the corresponding efficiency value?

****Solution Phase:****

- 1) The chart indicates that ****Nuclear**** has the highest energy efficiency.
- 2) Its efficiency is ****45.77%****, which corresponds to option ****C****.

****Final Answer:** C**

Qwen2.5-VL-3B (Base) : d

Qwen2.5-VL-3B (VisDoT) : Let's break down this problem.

****Question Phase:****


1. Which energy source has the highest energy efficiency according to the chart?
2. What is the exact value of its energy efficiency?

****Solution Phase:****

- 1) The chart shows that ****Nuclear**** (blue) has the highest bar.
- 2) Its energy efficiency is approximately ****48.77%****.

****Final Answer:** C**

Figure 13: ChartQAPro Multi Choice task



Question : Is the value at the highest point in the middle of the chart more than 1.5 percentage points greater than the value at the leftmost position ?

Gemma3-4B (Base) : Here's the analysis: ****Leftmost value:**** 3.8% (2012), ****Highest value:**** 5.3% (2010), ****Difference:**** $5.3\% - 3.8\% = 1.5\%$, ****Answer:**** **Yes**, the value at the highest point (5.3%) is 1.5 percentage points greater than the value at the leftmost position (3.8%).

Gemma3-4B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. What is the value at the highest point in the middle of the chart (2010)?
2. What is the value at the leftmost position (2005)?
3. Is the 2010 value more than 1.5 percentage points greater than the 2005 value?

****Solution Phase:****

- 1) The value at the middle high point in 2010 is ****4.1%****.
- 2) The value at the leftmost point in 2005 is ****4.2%****.
- 3) Subtracting the two values: ****4.1% - 4.2% = -0.1%****


Since the result is negative, the 2010 value is actually lower than the 2005 value.

****Final Answer:**** No

Figure 14: VisDoTQA Position task

Widespread concern about mobile phones' impact on children across 11 emerging economies surveyed

% of adults who say people should be very/somewhat/not concerned about ___ when using their mobile phones



Note: Percentages are 11-country medians.

Source: Mobile Technology and its Social Impact Survey 2018, Q19a-f.

"Mobile Connectivity in Emerging Economies"

PEW RESEARCH CENTER

Question : By how much does the combined total of the darkest blue bars exceed the combined total of the medium blue bars across all categories?

Gemma3-4B (Base) : Here's the breakdown of the analysis:

1. ****Identify the Darkest Blue Bars:**** These represent \"Very concerned\".

The total percentage for this category is **79%**.

2. ****Identify the Medium Blue Bars:**** These represent \"Somewhat concerned\".

The total percentage for this category is **17%**.

3. ****Calculate the Difference:**** 79% - 17% = **62%**

****Answer:**** The combined total of the darkest blue bars (very concerned) exceeds the combined total of the medium blue bars (somewhat concerned) by 62%.

Gemma3-4B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. What are the values of the dark blue bars (Very concerned) for each category?

2. What are the values of the medium blue bars (Somewhat concerned) for each category?

3. What is the combined total of the dark blue bars?

4. What is the combined total of the medium blue bars?

5. By how much does the total of the dark blue bars exceed that of the medium blue bars?

****Solution Phase:****

1) **Dark Blue Bars (Very concerned):**

- Losing the ability to communicate face-to-face: 48

- Harassment or bullying: 59

- Mobile phone addiction: 62

- Exposure to false or incorrect information: 64

- Identity theft: 66

- Children being exposed to harmful content: 79

2) **Medium Blue Bars (Somewhat concerned):**

- Losing the ability to communicate face-to-face: 26

- Harassment or bullying: 21

- Mobile phone addiction: 17

- Exposure to false or incorrect information: 25

- Identity theft: 17

- Children being exposed to harmful content: 14

3) **Total of dark blue bars:**

$48 + 59 + 62 + 64 + 66 + 79 = 378$

4) **Total of medium blue bars:**

$26 + 21 + 17 + 25 + 17 + 14 = 120$

5) **Difference:**

$378 - 120 = 258$

****Final Answer:**** The combined total of the dark blue bars exceeds the total of the medium blue bars by **258**.

Figure 15: VisDoTQA Pattern task

System Prompt: Perception-following Question Generation - Extract Prompt

You are an AI assistant that generates analytical questions based on numerical data in charts. Your task is to create **complex, meaningful questions** that require reasoning and analysis rather than simple data lookup.

Required Operations (use each at least twice across 20 questions):

- **Counting:** Identifying how many items meet a specific numerical condition (e.g., "How many regions exceed 20%?").
- **Addition & Subtraction:** Summing or finding differences between multiple values.
- **Multiplication & Division:** Calculating proportional changes or percentage comparisons.
- **Average & Median:** Comparing means and middle values across multiple categories or regions.
- **Ratio:** Determining how many times one value is larger or smaller than another.
- **Ranking:** Clearly identifying positions (highest, lowest, top 3, bottom 2, etc.).
- **Comparison:** Identifying higher/lower values or categories without performing explicit mathematical operations.


Important Guidelines

- Clearly state **at least two operations** per question.
- **All listed operations** must be used at least **twice across the 20 questions**.
- Avoid overly simple or single-operation arithmetic questions.
- Ensure each question encourages **analytical thinking and deeper interpretation** of chart data.

Example Questions

1. What is the combined percentage of the two highest-ranked categories? (*Addition, Ranking*)
2. By what percentage is the average of Category A higher than the median of Category B? (*Average, Median, Subtraction*)
3. How much higher is the 2nd highest value than the 4th highest value in the Sales column? (*Ratio, Ranking*)
4. Which three categories have the lowest values in the Good rating column? (*Ranking*)
5. Which country has the largest gap between negative and positive ratings? (*Subtraction, Ranking*)
6. How many regions have a value above the overall average in the Bachelor's degree category? (*Counting, Average, Comparison*)

Table 13: Perception-following Question Generation - Extract Prompt



Question : Between which two consecutive years did Fiji experience the largest increase in military expenditure as a share of GDP?

InternVL2.5-2B (Base) : **1978.0, 1980.0**

InternVL2.5-2B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. What are the military expenditure values for Fiji in consecutive years?
2. Between which two consecutive years is the largest change observed?

****Solution Phase:****


1) The chart shows the following values for Fiji's military expenditure as a share of GDP:

- 1970: ****0.19****
- 1972: ****0.26****
- 1974: ****0.25****
- 1976: ****0.41****
- 1978: ****0.89****
- 1980: ****0.87****

2) Comparing consecutive years, the largest difference occurs between ****1976 and 1978****, with a change of ****0.48**** (0.89 - 0.41).

****Final Answer:**** The largest increase in military expenditure as a share of GDP occurred between ****1976 and 1978****.

Figure 16: VisDoTQA Length task



Question : By how many years does the average expected years of schooling in Mongolia and Namibia exceed the expected years in Tajikistan?

Qwen2.5-VL-3B (Base) : To determine by how many years the average expected years of schooling in Mongolia and Namibia exceed the expected years in Tajikistan, we need to follow these steps:

1. Identify the expected years of schooling for each country from the chart:\n - **Mongolia: 11.8 years**\n - **Namibia: 11.6 years**\n - Tajikistan: 10.6 years
2. Calculate **the total expected years** of schooling for Mongolia and Namibia: $11.8 + 11.6 = 23.4$ years
3. Subtract the expected years of schooling in Tajikistan from the total: $23.4 - 10.6 = 12.8$ years

Therefore, the average expected years of schooling in Mongolia and Namibia exceed the expected years in Tajikistan by **12.8 years**.

Qwen2.5-VL-3B (VisDoT) : Let's break down this problem.

****Question Phase:****

1. What is the average expected years of schooling in Mongolia and Namibia?
2. By how many years does this average exceed the expected years in Tajikistan?

****Solution Phase:****

1) The expected years of schooling are:

- Mongolia: ****11.8****
- Namibia: ****11.6****

The average is:

$$(11.8 + 11.6) \div 2 = \mathbf{**11.7**}$$

2) Tajikistan's expected years of schooling is ****10.6****.

The difference is:

$$11.7 - 10.6 = \mathbf{**1.1**}$$

****Final Answer:**** The average expected years of schooling in Mongolia and Namibia exceeds that in Tajikistan by ****1.1 years****.

Figure 17: VisDoTQA Extract task

System Prompt: DoT Prompt

You are an AI assistant that analyzes chart images and constructs the reasoning process to derive answers when given question-answer pairs.

You must follow a structured approach by decomposing the question into sub-questions and solving them systematically.

Response Format:

—

Let's break down this problem.

Question Phase:

1. (Ask about element 1)
2. (Ask about element 2)
3. ...

Solution Phase:

—

- 1) (Answer sub-question 1 with reasoning and answer)
- 2) (Answer sub-question 2 with reasoning and answer)
- 3) ...

Guidelines:

—

- Question Phase:

- Identify the necessary elements from the question to solve the problem.
- Generate sub-questions **strictly following the format**:
 1. (Ask about element 1)
 2. (Ask about element 2)
 3. ...
- The number of sub-questions must be **exactly reflected in the Solution Phase**.

- Solution Phase:

- Ensure that every sub-question has a corresponding answer in the exact order.
 - Provide answers **strictly following the format**:
 - 1) (Answer sub-question 1 with reasoning and answer)
 - 2) (Answer sub-question 2 with reasoning and answer)
 - 3) ...
 - Responses must be **clear, precise, and follow the structured format without skipping or merging answers**.
 - **Do not include the word "table" in any response.**
-

Table 14: DoT Prompt

Category	Prompt Template
ChartQAPRO (fact checking)	You are a meticulous chart-analysis assistant. Based on the chart image, decide whether the following statement is true or false. Your final answer must be either 'True' or 'False'.
ChartQAPRO (multi choice)	You are a meticulous chart-analysis assistant. Read the chart image and select the most appropriate answer from the multiple-choice options in the question, returning only the final answer as a single letter: A, B, C, or D.
ChartQAPRO (hypothetical, factoid)	You are a meticulous chart-analysis assistant. Analyze the chart image and answer the following question.
ChartQAPRO (conversational)	You are a meticulous chart-analysis assistant. Your goal is to read the multi-turn conversation carefully and provide the answer to the final question. Conversation: {conversation_text} Question: {final_question}
ChartQA, VisDoT	You are a meticulous chart-analysis assistant. Analyze the chart image and answer the following question.

Table 15: Model Inference Prompt




Figure 18: Example of position-based reasoning.




Figure 19: Example of length-based reasoning.




Figure 20: Example of pattern-based reasoning.




Figure 21: Example of extractive reasoning.