

# VProChart: Answering Chart Question Through Visual Perception Alignment Agent and Programmatic Solution Reasoning

Muye Huang<sup>1,2</sup>, Lingling Zhang<sup>1,2\*</sup>, Lai Han<sup>1,2</sup>, Wenjun Wu<sup>1,3</sup>, Xinyu Zhang<sup>1,3</sup>, Jun Liu<sup>1,2</sup>

<sup>1</sup>School of Computer Science and Technology, Xi'an Jiaotong University

<sup>2</sup>MOE KLINNS Lab, Xi'an Jiaotong University

<sup>3</sup>Shaanxi Province Key Laboratory of Big Data Knowledge Engineering  
{huangmuye, hanlai, nickjun98, zhang1393869716}@stu.xjtu.edu.cn,  
{zhanglling, liukeen}@xjtu.edu.cn

## Abstract

Charts are widely used for data visualization across various fields, including education, research, and business. Chart Question Answering (CQA) is an emerging task focused on the automatic interpretation and reasoning of data presented in charts. However, chart images are inherently difficult to interpret, and chart-related questions often involve complex logical and numerical reasoning, which hinders the performance of existing models. This paper introduces VProChart, a novel framework designed to address these challenges in CQA by integrating a lightweight Visual Perception Alignment Agent (VPAgent) and a Programmatic Solution Reasoning approach. VPAgent aligns and models chart elements based on principles of human visual perception, enhancing the understanding of chart context. The Programmatic Solution Reasoning approach leverages large language models (LLMs) to transform natural language reasoning questions into structured solution programs, facilitating precise numerical and logical reasoning. Extensive experiments on benchmark datasets such as ChartQA and PlotQA demonstrate that VProChart significantly outperforms existing methods, highlighting its ability to understand and reason with charts.

**Homepage** — <https://github.com/MuyeHuang/VproChart>

## 1 Introduction

Chart is a widely used format for data visualization, prevalent in scientific papers and business reports. Chart Question Answering (CQA) aims to answer questions based on chart contexts, achieving the automatic analysis of data trends and the automatic generation of data reports. In recent years, the CQA task has garnered increasing attention in light of advancements in multi-modal understanding and reasoning techniques. (Kafle et al. 2018; Methani et al. 2020)

The CQA task naturally draws comparisons with Visual Question Answering (VQA), which involves answering questions based on images. Despite significant advances in VQA (Li et al. 2023; Wang et al. 2023), CQA remains challenging due to the complex understanding of diagram contexts and complex reasoning requirements. Firstly, chart images are inherently difficult to interpret. For example, in

\*Corresponding author.

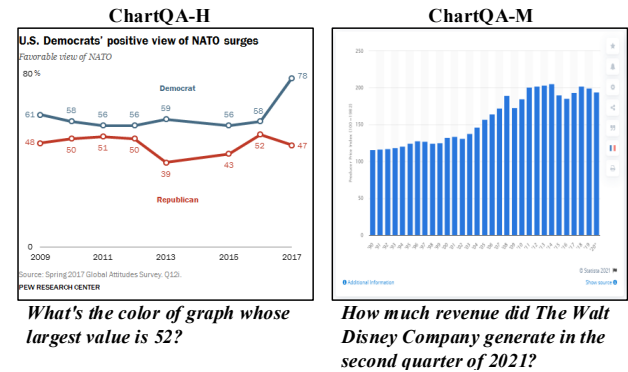


Figure 1: Comparison of the two subsets of ChartQA, ChartQA-H and ChartQA-M.

the left chart of Figure 1, the line chart does not have a strict legend, it only provides two labels corresponding to colors, which might potentially cause confusion. The elements and organizational methods in charts are not derived from nature but depend on the author's intentions. Different authors may use various elements or organizational methods for clarity or aesthetic appeal, complicating the understanding of chart images through common presentation rules. Secondly, questions in CQA often require complex numerical or logical reasoning. For example, in the left chart of Figure 1, "Is the sum of the heights of the two highest red line greater than the value of the two highest blue line?" necessitates numerical calculations and logical comparisons. These types of reasoning are particularly challenging for current research methods. However, existing approaches lack significant improvement in both of these aforementioned aspects.

On the one hand, the currently dominant strategy includes models like Pix2Struct (Lee et al. 2023) and UniChart (Masry et al. 2023), which are end-to-end visual models, using extensive synthetic chart pretraining to enhance their chart understanding capabilities. This method has made progress in understanding synthetic charts but is still limited by the performance of external tools and synthetic data. On the other hand, MatCha (Liu et al. 2023) and ChartInstruct (Masry et al. 2024a) have successfully enhanced the reasoning capabilities of models on CQA tasks through

logical reasoning computation datasets or instruction fine-tuning. However, these reasoning approaches are still based on neural networks, making it difficult to achieve further progress in numerical reasoning. Overall, existing research lacks improvements for the two aforementioned challenges: in chart understanding, current methods rely on pretraining with large-scale chart-question pairs, while ignoring explicit modeling of relationships among chart elements; in numerical and logical reasoning, they rely on pretrained neural networks, which are not adept at handling complex logical and numerical reasoning.

To address the above two issues, we propose VProChart, an Agent-LLMs hybrid framework designed for the challenging CQA task. The VProChart method consists of two modules: the lightweight visual perception alignment agent and the programmatic solution reasoning approach. The Visual Perception Alignment agent (VPAgent) achieves chart understanding and basic chart question answering by principles of human visual perception. The programmatic solution reasoning approach converts questions into solution programs, which guide the reasoning process. During this process, the solution programs invokes the VPAgent to obtain intermediate variable values specified in the solution programs. Additionally, we perform a series of comparative experiments to demonstrate the superiority of VProChart.

Our main contributions can be summarized into four folds:

- We propose the novel VProChart method to solve the challenges in understanding the context of the chart and the complex reasoning demands of the CQA task.
- We explicitly model and align the relationships between chart elements based on human visual perception principle, and infer answers within the modeled relationships using a question-driven reasoning module.
- We propose a programmatic solution reasoning framework, a flexible framework that can be applied to all types of questions, accurately inferring answers through flexible solution approaches and a precise Python-style reasoning engine.
- We conduct extensive experiments on different datasets such as ChartQA and PlotQA to validate the superiority of VProChart. The results show significant improvements compared to various popular competitors across different scales.

## 2 Related Work

The chart question answering task answers textual questions by understanding visual context. Kahou et al. (Kahou et al. 2018) pioneered a synthetic Chart dataset FigureQA, which includes automatically generated chart images and a small number of template-generated true/false judgment questions. Kafle et al. (Kafle et al. 2018) proposed another more comprehensive synthetic dataset, DVQA. They also introduced two models named MOM and SANDY to better handle CQA problems, and they subsequently introduced the PReFIL (Kafle et al. 2020) model, achieving state-of-the-art results on FigureQA and DVQA.

However, the aforementioned datasets lack the complex reasoning problems present in real-world CQA. Therefore, PlotQA (Methani et al. 2020) introduced a large-scale dataset, containing complex reasoning and numerical computation tasks. It comprises real-world data along with diverse synthetic charts. Furthermore, Masry et al. (Masry et al. 2022) proposed the novel CQA dataset ChartQA, where the human subset consists of real-world images and manually curated questions, bringing the CQA task closer to real-world scenarios. CRCT (Levy, Ben-Ari, and Lischinski 2022) achieved satisfactory results on the PlotQA dataset by utilising a unified regression-classification head. Lee et al. (Lee et al. 2023) introduced a pure visual model Pix2Struct based on Vision Transformer (Dosovitskiy et al. 2021), achieving state-of-the-art performance on various CQA datasets through large-scale pre-training on screenshots without the need for external OCR. Building upon this, MatCha (Liu et al. 2023) using mathematical reasoning tasks during pre-training, resulting in significant improvements on ChartQA and PlotQA. Masry et al. (Masry et al. 2023) proposed UniChart, which achieved state-of-the-art performance on the ChartQA task through multi-task multi-stage pre-training. In light of the rapid advancement of LLMs, ChartLlama (Han et al. 2023) utilised Llama with an added visual branch, achieving good performance through large-scale pre-training and offering new directions for researchers. Moreover, ChartInstruct (Masry et al. 2024a), ChartAst (Meng et al. 2024) and TinyChart (Zhang et al. 2024), an instructive approach to the CQA domain, which achieved new advancements in CQA tasks through the utilisation of LLMs.

The aforementioned works either rely on external detectors and data annotations or heavily on large-scale synthetic data driving, neglecting the exploration of structural relationships among chart elements and failing to achieve precise reasoning, limiting their understanding of charts and the reasoning capability for questions. In contrast, our work focuses on explicit modeling of chart structural relationships and precise question reasoning, achieving impressive results across multiple datasets.

## 3 Method

In this section, we introduce VProChart method, which comprises a lightweight visual perception alignment agent and programmatic solution reasoning approach. The programmatic solution reasoning approach transforms the reasoning problem into a programmatic solution and invokes the visual perception alignment agent to understand the chart context for reasoning execution.

### 3.1 Visual Perception Alignment Agent

The Visual Perception Alignment Agent (VPAgent) is designed for chart question answering. Figure2 illustrates the framework of the VPAgent, which includes four modules: 1) Chart & question encoder module to extract sequences from charts and questions; 2) Visual perception align module to explicitly model the relationships between elements; 3) Q-driven chart reasoning module to infer and provide answers

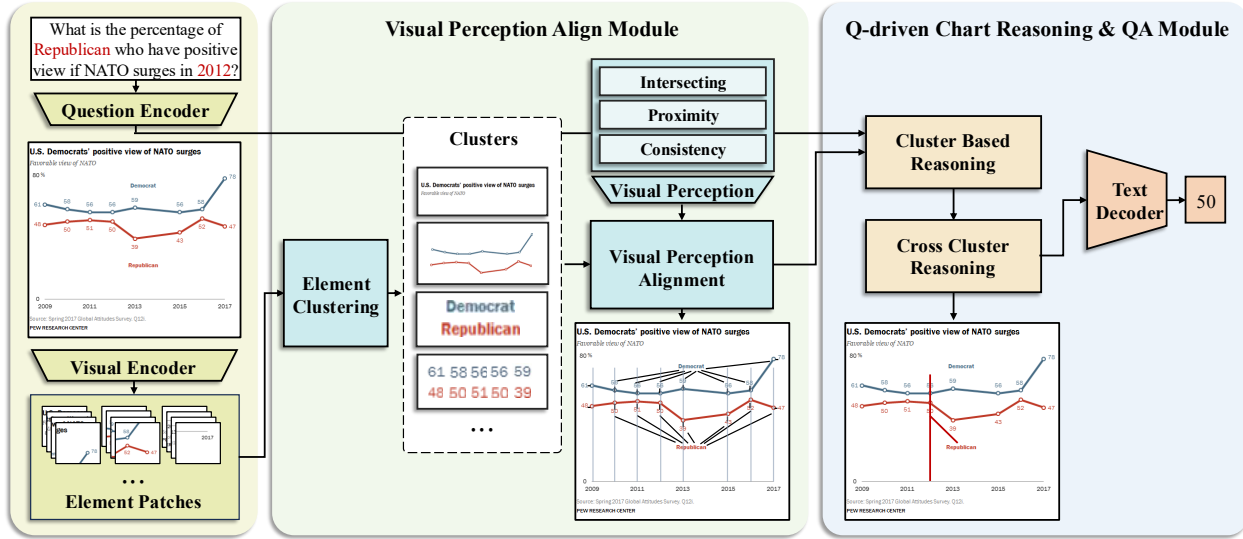


Figure 2: The overview of the proposed VPAgent. The proposed VPAgent comprises four modules: Chart & Question Encoder Module (left), Visual Perception Align Module (center), Q-driven Chart Reasoning and Question Answering Module (right).

driven by the question; 4) Question Answering Module to produce the answer. The details of these four modules are introduced as follows:

**1) Chart & Question Encoder Module.** The Chart & Question Encoder Module converts the chart and question into sequential representations. This encoder is designed to recognize both textual and visual elements present in  $C_v$  in the chart, and to represent the question in a sequential format that encapsulates the core essence  $Q_{global}$  of the question.

**Chart Encoder Module.** The encoding of chart images plays a critical role in the CQA process, as it requires the extraction of both visual and textual elements from the charts. To address the need for a lightweight and structurally complex encoder, we employ Swin-Encoder (Liu et al. 2021) as the chart encoder. Swin-Encoder is a lightweight encoder capable of achieving high-resolution encoding at a relatively low computational cost, enabling the simultaneous extraction of visual and textual information from images without the need for external OCR.

$$C_v = \{X_1, X_2, \dots, X_n\} \quad (1)$$

where  $C_v$  is a sequence set containing chart information, and  $X_k$  represents the embedding of the  $k$ -th patch.

**Question Encoder Module.** The Question Encoder Module encodes questions into sequences and extracts a global representation of each question. We employ RoBERTa (Liu et al. 2019), renowned for its state-of-the-art performance in natural language processing tasks, as the encoder. Given a question  $Q$ , this procedure represents the question as a sequence:

$$Q_t = \{H_0, H_1, H_2, \dots, H_m\} \quad (2)$$

where  $Q_t$  is the sequential representation of question  $Q$ , and  $Q_{global}$  is the global embedding of the question.

**2) Visual Perception Align Module.** In the image sequence set  $C_v$ , each patch contains elements of the chart. This module clusters elements of the same type within the image sequence. It then models the structural relationships among these elements by aligning elements from different clusters based on human visual perception principles. Finally, this module represents the visual perception alignment information as a matrix  $V$ .

**Element Clustering.** Chart images come in various forms (e.g., pie charts, line charts, bar charts), each with distinct element names and features. To address this diversity of elements, we employ agglomerative hierarchical clustering, a flexible and noise-robust approach. This method operates at the patch level within chart images, organizing the differentiated elements into clusters:

$$G = \{G_0, G_1, \dots, G_{k-1}\}, \quad |G| = k \quad (3)$$

$$C_v = \bigcup_{i=0}^{k-1} G_i, \quad G_i \cap G_j = \emptyset \quad \forall i \neq j \quad (4)$$

where  $k$  denotes the number of clusters,  $G$  is the set of all clusters, and  $C_v$  is the set of image embedding sequences.

**Visual Perception Alignment.** In this module, we align elements from different clusters based on visual perception principles: **a) Crosshair Intersecting Principle**, aligning elements based on their horizontal and vertical relationships on the original chart image, e.g., aligning the axes with the lines or bars. This is because charts always adhere to the Crosshair Intersecting alignment relationship for the  $x$  and  $y$  axes. **b) Spatial Proximity Principle**, alignment is determined by the spatial proximity of elements on the original chart, e.g., aligning bars with their corresponding numerical labels. This is because elements that are spatially proximate within the chart are more likely to be related. **c) Color**

**& Semantic Consistency Principle**, alignment is based on the color and semantic similarity of elements, e.g., aligning green legend labels with green bars. This is because elements of the same color in the chart commonly represent the same attribute.

These three principles of alignment relationships can be respectively represented as **I**, **P**, and **S**. For **I**, we establish relationships with patches that are in the same row or column as patch  $X_i$ , and store the collection of these relationships within it. For **P**, we establish relationships with patches that are within a straight-line distance of 3 from patch  $X_i$ , and store the collection of these relationships within it. For **S**, we calculate the normalized semantic similarity with patch  $X_i$  and store these relationships within it.

The primary alignment method varies among different types of charts. For instance, pie charts predominantly employ spatial proximity, while line graphs primarily utilize crosshair intersecting. We utilize the global pooling representation  $\mathbf{X}_p$  of the chart images to obtain the weights  $\mathbf{W}_c$  of the three alignment through a MLP:

$$\mathbf{W}_c = \text{softmax}(\mathbf{W}_2 \cdot \sigma(\mathbf{W}_1 \cdot \mathbf{X}_p + \mathbf{b}_1) + \mathbf{b}_2) \quad (5)$$

However, not every pair of clusters in the chart needs to be aligned. The necessity of alignment between clusters is determined by computing the interactions among the global representations of each cluster. This ensures that the alignment process is effective, focusing only on the most relevant cluster interactions. Let  $\mathbf{G}$  be the matrix of global representations of clusters:

$$\mathbf{G} = [\mathbf{g}_0, \mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_{k-1}] \quad (6)$$

The interaction matrix  $\mathbf{W}$ , which denotes the necessity scores of alignment between clusters, is computed as:

$$\mathbf{W} = \text{softmax}(\mathbf{G} \cdot \mathbf{G}^T) \quad (7)$$

By combining the weights  $\mathbf{W}_c$  and  $\mathbf{W}$ , we integrate the alignment relationships between clusters into a visual criterion alignment matrix  $\mathbf{V}$ :

$$\mathbf{V} = \mathbf{W}_c \odot (\mathbf{W} \odot (\mathbf{I} + \mathbf{P} + \mathbf{S})) \quad (8)$$

Here,  $\odot$  denotes the element-wise multiplication. Overall, the visual perception alignment module explicitly provides the alignment information  $\mathbf{V}$  for chart image elements by integrating visual perception principles and multi-weight alignment information.

**3) Q-driven Chart Reasoning QA Module.** Given a chart  $\mathbf{C}_v$ , this module mines implicit information within the graph and marks answers under the guidance of the problem-global representation  $\mathbf{H}_0$  and the chart alignment information  $\mathbf{V}$ . To achieve this, we designed the intra-cluster reasoning module and the cross-cluster reasoning module to handle intra-cluster and inter-cluster reasoning tasks respectively.

**Intra-Cluster Reasoning.** In charts, certain information such as omitted interval labels is implicitly provided. This module deeply mines the implicit information within clusters by using a restricted-attention mechanism. The mechanism confines the attention scope to the range of element

clusters  $\mathbf{G}_i$ . It is followed by feed-forward layers, skip connections, and normalization. We stack this attention module for  $\mathbf{N}$  layers to enhance information extraction.

$$\mathbf{C}_r = \sum_{i=0}^k \left[ \text{Softmax} \left( \frac{\mathbf{W}_Q \mathbf{C}_v [\mathbf{W}_K \mathbf{G}_i]^T}{\sqrt{d_k}} \right) \mathbf{W}_V \mathbf{G}_i \right]^N \quad (9)$$

Here,  $\mathbf{C}_r$  represents the output image sequence set of this module.

**Cross-Cluster Reasoning.** This module utilizes the alignment matrix  $\mathbf{V}$  and the global embedding of the question  $\mathbf{Q}_{\text{global}}$  to annotate elements in the alignment information  $\mathbf{V}$  that are strongly associated with the answer. Specifically, it computes the similarity scores between image elements in  $\mathbf{C}_r$  and the global representation  $\mathbf{Q}_{\text{global}}$  of the question. The top-k elements most relevant to the query are then selected and annotated in the alignment matrix  $\mathbf{V}$ :

$$\mathbf{V}_r = \text{Annotate}(\text{softmax}(\mathbf{C}_r \mathbf{Q}_{\text{global}}^T), \mathbf{V}) \quad (10)$$

Here,  $\mathbf{V}_r$  represents the reasoned alignment matrix with the most relevant elements annotated. As the alignment matrix  $\mathbf{V}$  only contains alignment relationships between elements from different clusters, marking one element results in automatically marking elements from other clusters that are aligned with it. Consequently, the reasoned chart  $\mathbf{C}_r$  obtained from the previous steps is combined with the reasoned alignment information  $\mathbf{V}_r$ , yielding the output  $\mathbf{O}_c$ . This combination is achieved through a normalization process, which ensures the stability and consistency of the output:

$$\mathbf{O}_c = \frac{\mathbf{C}_r + \mathbf{V}_r - \mu}{\sqrt{\sigma^2 + \epsilon}} \gamma + \beta \quad (11)$$

Here,  $\mu$  and  $\sigma^2$  represent the mean and variance of  $\mathbf{C}_r + \mathbf{V}_r$ , respectively, while  $\epsilon$  is a small constant for numerical stability. The parameters  $\gamma$  and  $\beta$  are learnable affine transformation parameters.

**4) Question Answering Module.** For a given question  $\mathbf{Q}$ , the answer can be either a floating-point number or textual content. Previous works, such as T5 (Raffel et al. 2020), have demonstrated that generative decoders can effectively handle both types of responses simultaneously. Consequently, we employ the mBART (Liu et al. 2020) decoder as the text decoder. The text prompt is provided as input to the decoder, which generates the answer based on the prompt context. The loss function for this process is defined as follows:

$$\mathcal{L} = - \sum_{t=1}^T \log p(y_t | y_{<t}, \mathbf{O}_c, \mathbf{Q}) \quad (12)$$

where  $y_t$  represents the target token at time step  $t$ ,  $y_{<t}$  denotes the sequence of previous target tokens, and  $p(y_t | y_{<t}, \mathbf{O}_c, \mathbf{Q})$  is the predicted probability of  $y_t$  given  $y_{<t}$ ,  $\mathbf{O}_c$ , and  $\mathbf{Q}$ . This loss function minimizes the negative log-likelihood of the predicted tokens.

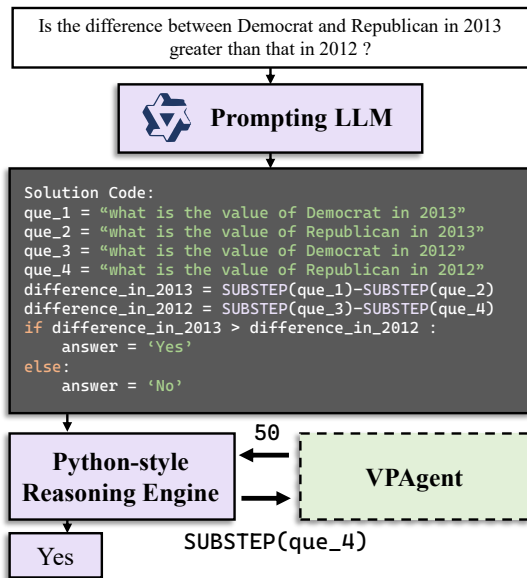


Figure 3: An example of Programmatic Solution Reasoning.

### 3.2 Programmatic Solution Reasoning

CQA necessitates advanced reasoning capabilities. To address this challenge, we propose the Programmatic Solution Reasoning approach, which leverages LLMs to integrate the precise reasoning of programs with the language understanding of LLMs. As depicted in Figure 2, the framework consists of two distinct modules: 1) Prompting LLMs to parse natural language questions into structured pseudo-code solution in Python style. 2) Python-style Reasoning Engine to interpret the solution generated by the LLMs and output answers through programming.

#### 1) Prompting LLMs.

Prompting LLMs leverage the complex question understanding and code generation capabilities of LLMs to transform intricate textual questions into solution program (Xu et al. 2023). LLMs do not need to compute or infer answers, instead, they express the process of solving the problem in the form of solution program (Chen et al. 2023a). In particular, we prompt the LLMs to act as programming experts, solving problems programmatically by invoking two kinds of operators: 1) *ASK*( $\cdot$ ). This operator takes the complete question as input and returns the answer directly, used for questions that do not require reasoning or computation. For instance, the question “Which color indicates 65+ years?” requires no reasoning. 2) *SUBSTEP*( $\cdot$ ). This operator takes in sub-questions and returns answers to these sub-questions, used for complex questions requiring decomposition.

#### 2) Python-style Reasoning Engine.

Given any question  $Q$  and its corresponding solution  $P$ , this module parses the solution program  $P$  and employs the VPAgent to obtain values for critical variables within the program. The flexibility of the programming design endows the reasoning engine with broad applicability to various types of questions, ensuring robust performance across

diverse query scenarios. For example, as shown in Figure 3, the key variable *difference\_in\_2012* in the solution program requires calling the VPAgent twice to obtain the answer. The Reasoning Engine passes the variable “what is the value of Republican in 2012” from the SUBSTEP to the VPAgent, and returns the output of the Agent, “50.” Finally, the Reasoning Engine uses a judgment statement to obtain the answer output of the solution program. This process ensures the accuracy of the reasoning.

## 4 Experiments

### 4.1 Experimental Setting

**Datasets.** We conducted extensive experiments on three prominent datasets: PlotQA (Methani et al. 2020), DVQA (Kafle et al. 2018), and ChartQA (Masry et al. 2022). DVQA is a fully synthetic dataset focused on chart-specific questions. PlotQA, a large synthetic dataset, comprises two subsets, D1 and D2, incorporating both real-world data and automatically generated charts, with numerous numerical and logical reasoning questions. ChartQA includes two benchmarks for challenging CQA tasks: ChartQA-H and ChartQA-M. ChartQA-H features manually curated chart images and question-answer pairs, while ChartQA-M consists of machine-generated images and questions. As illustrated in Figure 1, ChartQA-H poses a more formidable challenge due to its extensive collection of manually curated charts and complex multi-step reasoning questions.

**Competitors.** We evaluate our model against both fully-supervised models and large-language-model-based baselines. *Fully-supervised models* include T5 (Raffel et al. 2020), PReFIL (Kafle et al. 2020), VisionTapas (Masry et al. 2022), Pix2Struct (Lee et al. 2023), ChartReader (Cheng, Dai, and Hauptmann 2023), CRCT (Levy, Ben-Ari, and Lischinski 2022), MatCha (Liu et al. 2023), and UniChart (Masry et al. 2023), all of which are trained on CQA benchmarks. *Large-language-model-based models* such as PaLI-17B (res.588) (Chen et al. 2023b), Qwen-VL (Bai et al. 2023), ChartLLama (Han et al. 2023), OneChart (Chen et al. 2024), ChartGemma (Masry et al. 2024b) and ChartInstruct (Masry et al. 2024a) utilize end-to-end LLMs integrated with a visual component.

**Settings.** In our proposed VPAgent, we use the Donut weights (Kim et al. 2022) for initialization. The training process consists of two stages: firstly, pre-training 200k steps on the UniChart pretraining corpus (Masry et al. 2023), with an initial learning rate of  $1e-4$ ; secondly, fine-tuning 5 epoch on the downstream dataset, with an initial learning rate of  $5e-5$ . Notably, due to the substantial computational cost, we trained the model only on PlotQA-D1 and tested it on PlotQA-D2. In our programmatic solution reasoning framework, we employ *Qwen2-7B-Instruct* (Yang et al. 2024) as the base LLM. All experiments were conducted using two Tesla A100 GPUs.

### 4.2 Performance Comparison

In Table 1 and 2, we present a comprehensive summary of VProChart’s performance on three popular datasets. From these results, we make the following three observations.

Model	Size	ChartQA			PlotQA		DVQA		
		Human	Augment	Avg	PlotQA-D1	PlotQA-D2	Familiar	Novel	Avg
T5	223M	25.1	41.0	33.1	72.6	56.2	89.0	76.9	83.0
PReFIL	-	-	-	-	57.9	10.4	80.9	80.0	80.5
VisionTapas	223M	29.6	61.4	45.5	65.3	42.5	94.4	95.5	94.9
CRCT	-	-	-	-	76.9	34.4	-	-	82.1
Pix2Struct	300M	30.5	81.6	56.0	73.2	71.9	-	-	-
ChartReader	-	-	-	52.6	78.1	59.3	<u>95.4</u>	<u>96.5</u>	<u>96.0</u>
MatCha	300M	38.2	<b>90.2</b>	64.2	92.3	<b>90.7</b>	-	-	-
UniChart	260M	43.9	87.8	65.8	28.6	21.7	-	-	-
VPAgent(Ours)	311M	<u>48.1</u>	<u>87.4</u>	<u>67.8</u>	<u>92.6</u>	<u>78.1</u>	<b>98.5</b>	<b>97.1</b>	<b>97.8</b>
VProChart(Ours)	7B+0.3B	<b>57.4</b>	<u>88.0</u>	<b>72.7</b>	<b>94.7</b>	<u>83.5</u>	-	-	-

Table 1: Comparison results on ChartQA, PlotQA, and DVQA Datasets

Model	Size	ChartQA		
		Human	Aug	Avg
PaLI(res.588)	17B	30.4	64.9	47.6
Qwen-VL	9.6B	44.3	78.9	61.6
ChartLlama	13B	48.9	<u>90.4</u>	69.7
ChartInstruct-E	7B	45.5	87.8	66.7
ChartInstruct-P	3B+0.3B	<u>50.1</u>	<b>93.8</b>	<u>71.9</u>
OneChart	13B+0.2B	49.1	85.3	67.2
VProChart(Ours)	7B+0.3B	<b>57.4</b>	88.0	<b>72.7</b>

Table 2: Comparison VLM results on ChartQA Dataset

1) Our VPAgent and VProChart method consistently outperform other competitors across three datasets in terms of overall performance. When comparing our VProChart method with similarly sized large models, the performance is nearly 5.0% higher compared to the best competitors on ChartQA. Additionally, when the small VPAgent is used for basic chart question-answering tasks and compared with similarly sized fully supervised models, its performance is higher than the best competitor UniChart by 2%. This indicates that our proposed VProChart method is advantageous for CQA tasks.

2) Compared to the best competitor ChartInstruct-P, our VProChart method achieves a performance improvement of over 7% on “human” set. Even without utilizing additional training data, the VPAgent surpasses the best competitor, UniChart, by more than 4% on the human subset. It’s worth noting that the gap between the small VPAgent and ChartLlama on the human subset is less than 1%. This illustrates that the proposed VProChart approach exhibits strong generalization and is suitable for real-world CQA tasks.

3) With all methods, the performance scores worsen when the benchmark changes from augment to human, or from PlotQA-D1 to D2. This is reasonable because the charts in the human subset are manually curated, and the questions entail multi-step complex reasoning. Although PlotQA-D2 lacks manually curated charts, most questions still require multi-step reasoning.

Method	ChartQA		
	Human	Augment	Avg
w/o VP alignment	44.1	87.0	65.6
w/o Q-driven chart reasoning	44.8	87.4	66.1
w/o intra-cluster reasoning	46.0	87.2	66.6
w/o cross-cluster reasoning	46.6	87.4	67.0
VPAgent	<b>48.1</b>	<b>87.4</b>	<b>67.8</b>

Table 3: Ablation study on VPAgent.

### 4.3 Analysis on VPAgent

**VPAgent Ablation.** To validate the contribution of the VP alignment module and the Q-driven reasoning module within VPAgent, we designed detailed ablation experiments for VPAgent, as shown in Table 3. Further experiments, including the zeroshot study and architecture study, can be found in the appendix. We use “w/o” to indicate the absence of a particular module. It is noteworthy that since VP alignment is a prerequisite module for Q-driven chart reasoning, in the “w/o VP alignment” experiment, we substitute the alignment matrix  $\mathbf{V}$  with Gaussian noise. We can make the following three observations:

1) The performance of models missing any component is significantly lower than VPAgent, demonstrating the effectiveness of our proposed VPAgent for general chart question answering.

2) Models with the Q-driven chart reasoning module outperform those without it, highlighting the overall effectiveness of the Q-driven chart reasoning module. However, the model without VP alignment exhibits the lowest performance, indicating the importance of incorporating visual perception into the model.

3) For different versions of the ablation models, the “human” performance drops much more dramatically than “Augment”. This is reasonable because “Augment” involves machine-generated images and questions, where answering them does not require the same level of visual complexity as “human” questions.

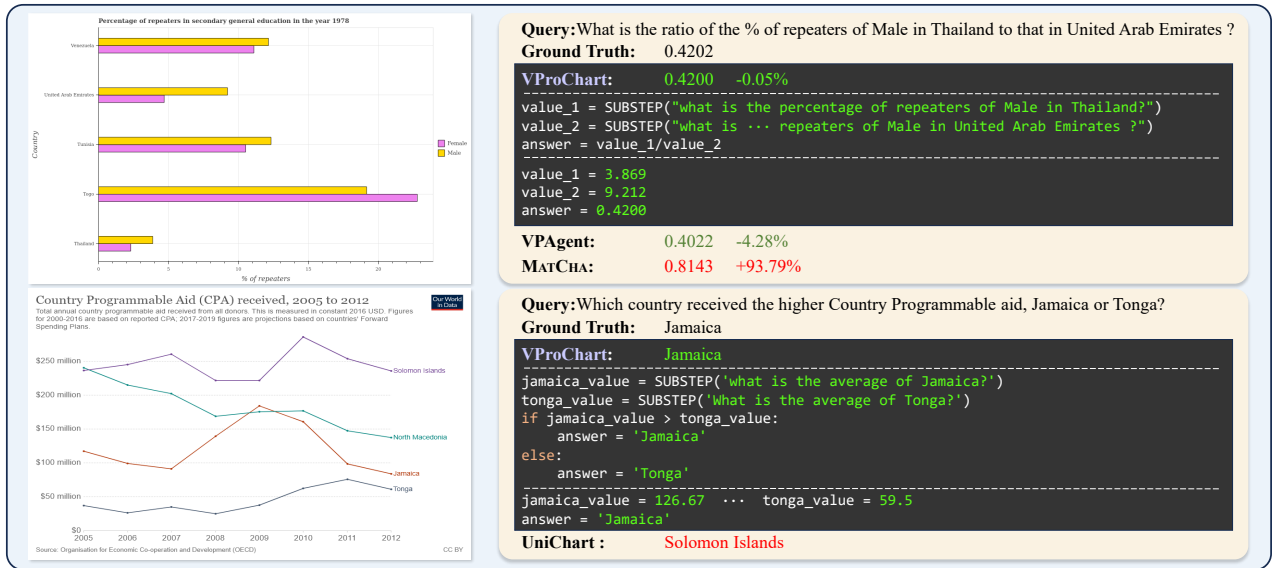


Figure 4: Two examples of VProChart. Above and below illustrate two examples from the PlotQA and ChartQA datasets, respectively.

Agent Method	ChartQA-Human	
	w/o Programmatic Solution	w/ Programmatic Solution
ChartInstruct-E	22.5	28.6 $\uparrow$ <b>6.1</b>
ChartInstruct-P	37.4	43.2 $\uparrow$ <b>5.8</b>
ChartGemma	28.0	34.9 $\uparrow$ <b>6.9</b>
UniChart	43.9	52.3 $\uparrow$ <b>8.4</b>
VPAgent	48.1	57.4 $\uparrow$ <b>9.3</b>

Table 4: Ablation study on Programmatic Solution Reasoning. It is noteworthy that all experiments were reproduced locally.

#### 4.4 Analysis on Programmatic Solution

**Programmatic Solution Ablation.** In this section, we utilize various chart expert models as agents to investigate the effectiveness of Programmatic Solution Reasoning. We use “w/o Programmatic Solution” to indicate the performance of the original agent and “w/ Programmatic Solution” to represent the performance of ChartQA under the guidance of the Programmatic Solution. Considering that the “Augmented” set almost lacks reasoning issues, our experiments focus on the ChartQA-Human set, which is rich in reasoning problems. The experimental results, as shown in Table 4, indicate that all agents exhibit significant performance improvements under the guidance of the Programmatic Solution. It demonstrates that Programmatic Solution Reasoning is a highly effective method for chart question answering.

#### 4.5 Case Study

To further investigate the capabilities of VProChart, we selected sample instances from the PlotQA and ChartQA datasets. Figure 4 illustrates an example from each datasets. In comparison to MatCha or UniChart, our VProChart yields predictions that are closer to the ground truth, even with-

out a Programmatic Solution. This observation underscores the effectiveness of VProChart. Further illustrative examples can be found in the Appendix.

## 5 Conclusion

In this paper, we propose the VProChart, a framework comprising a VPAgent guided by human visual principles and Programmatic Solution Reasoning for precise inference, which presents a novel approach to the ChartQA task. We conducted comprehensive experiments on ChartQA, PlotQA, and DVQA, demonstrating the effectiveness of VProChart, and showing that human visual perception can significantly enhance the model’s capability in understanding schematic diagrams. In the future, we will continue to explore the potential of integrating human-like thinking with programmatic reasoning.

## Acknowledgments

This work was supported by National Key Research and Development Program of China (2022YFC3303600), the Key Research and Development Project in Shaanxi Province No. 2022GXLH-01-03, National Natural Science Foundation of China (62137002, 62192781, 62293553, 61937001, 62250066, 62306229, and 62106190), “LENOVO-XJTU” Intelligent Industry Joint Laboratory Project, Foundation of Key National Defense Science and Technology Laboratory (6142101210201), Natural Science Basic Research Program of Shaanxi (2023-JC-YB-593), the Youth Innovation Team of Shaanxi Universities, the Shaanxi Provincial Social Science Foundation Project (2024P041).

## References

Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A Versatile Vision-

- Language Model for Understanding, Localization, Text Reading, and Beyond. *arXiv preprint arXiv:2308.12966*.
- Chen, J.; Kong, L.; Wei, H.; Liu, C.; Ge, Z.; Zhao, L.; Sun, J.; Han, C.; and Zhang, X. 2024. OneChart: Purify the Chart Structural Extraction via One Auxiliary Token. *arXiv preprint arXiv:2404.09987*.
- Chen, W.; Ma, X.; Wang, X.; and Cohen, W. W. 2023a. Program of Thoughts Prompting: Disentangling Computation from Reasoning for Numerical Reasoning Tasks. *TMLR*, 2835–8856.
- Chen, X.; Wang, X.; Changpinyo, S.; Piergiovanni, A. J.; Padlewski, P.; Salz, D.; Goodman, S.; Grycner, A.; Mustafa, B.; Beyer, L.; Kolesnikov, A.; Puigcerver, J.; Ding, N.; Rong, K.; Akbari, H.; Mishra, G.; Xue, L.; Thapliyal, A. V.; Bradbury, J.; and Kuo, W. 2023b. PaLI: A Jointly-Scaled Multilingual Language-Image Model. In *ICLR*.
- Cheng, Z.; Dai, Q.; and Hauptmann, A. G. 2023. ChartReader: A Unified Framework for Chart Derendering and Comprehension without Heuristic Rules. In *ICCV*, 22145–22156.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Housley, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Han, Y.; Zhang, C.; Chen, X.; Yang, X.; Wang, Z.; Yu, G.; Fu, B.; and Zhang, H. 2023. ChartLlama: A Multimodal LLM for Chart Understanding and Generation. *arXiv preprint arXiv:2311.16483*.
- Kafle, K.; Price, B. L.; Cohen, S.; and Kanan, C. 2018. DVQA: Understanding Data Visualizations via Question Answering. In *CVPR*, 5648–5656.
- Kafle, K.; Shrestha, R.; Price, B. L.; Cohen, S.; and Kanan, C. 2020. Answering Questions about Data Visualizations using Efficient Bimodal Fusion. In *WACV*, 1487–1496.
- Kahou, S. E.; Michalski, V.; Atkinson, A.; Kádár, Á.; Trischler, A.; and Bengio, Y. 2018. FigureQA: An Annotated Figure Dataset for Visual Reasoning. In *ICLR*.
- Kim, G.; Hong, T.; Yim, M.; Nam, J.; Park, J.; Yim, J.; Hwang, W.; Yun, S.; Han, D.; and Park, S. 2022. OCR-Free Document Understanding Transformer. In *ECCV*, 13688: 498–517.
- Lee, K.; Joshi, M.; Turc, I. R.; Hu, H.; Liu, F.; Eisenschlos, J. M.; Khandelwal, U.; Shaw, P.; Chang, M.; and Toutanova, K. 2023. Pix2Struct: Screenshot Parsing as Pretraining for Visual Language Understanding. In *ICML*, 202: 18893–18912.
- Levy, M.; Ben-Ari, R.; and Lischinski, D. 2022. Classification-Regression for Chart Comprehension. In *ECCV*, 13696: 469–484.
- Li, H.; Huang, J.; Jin, P.; Song, G.; Wu, Q.; and Chen, J. 2023. Weakly-Supervised 3D Spatial Reasoning for Text-Based Visual Question Answering. *TIP*, 32: 3367–3382.
- Liu, F.; Piccinno, F.; Krichene, S.; Pang, C.; Lee, K.; Joshi, M.; Altun, Y.; Collier, N.; and Eisenschlos, J. M. 2023. MatCha: Enhancing Visual Language Pretraining with Math Reasoning and Chart Derendering. In *ACL*, 12756–12770.
- Liu, Y.; Gu, J.; Goyal, N.; Li, X.; Edunov, S.; Ghazvininejad, M.; Lewis, M.; and Zettlemoyer, L. 2020. Multilingual Denoising Pre-training for Neural Machine Translation. *TACL*, 8: 726–742.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *ICCV*, 9992–10002.
- Masry, A.; Kavehzadeh, P.; Long, D. X.; Hoque, E.; and Joty, S. 2023. UniChart: A Universal Vision-language Pre-trained Model for Chart Comprehension and Reasoning. In *EMNLP*, 14662–14684.
- Masry, A.; Long, D. X.; Tan, J. Q.; Joty, S. R.; and Hoque, E. 2022. ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning. In *Findings of ACL*, 2263–2279.
- Masry, A.; Shahmohammadi, M.; Parvez, M. R.; Hoque, E.; and Joty, S. 2024a. ChartInstruct: Instruction Tuning for Chart Comprehension and Reasoning. *arXiv preprint arXiv:2403.09028*.
- Masry, A.; Thakkar, M.; Bajaj, A.; Kartha, A.; Hoque, E.; and Joty, S. 2024b. ChartGemma: Visual Instruction-tuning for Chart Reasoning in the Wild. *arXiv:2407.04172*.
- Meng, F.; Shao, W.; Lu, Q.; Gao, P.; Zhang, K.; Qiao, Y.; and Luo, P. 2024. ChartAssistant: A Universal Chart Multimodal Language Model via Chart-to-Table Pre-training and Multitask Instruction Tuning. *arXiv preprint arXiv:2401.02384*.
- Methani, N.; Ganguly, P.; Khapra, M. M.; and Kumar, P. 2020. PlotQA: Reasoning over Scientific Plots. In *WACV*, 1516–1525.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *JMLR*, 21:140:1–140:67.
- Wang, Y.; Wei, B.; Liu, J.; Zhang, L.; Wang, J.; and Wang, Q. 2023. DisAVR: Disentangled Adaptive Visual Reasoning Network for Diagram Question Answering. *TIP*, 32:4812–4827.
- Xu, F.; Wu, Z.; Sun, Q.; Ren, S.; Yuan, F.; Yuan, S.; Lin, Q.; Qiao, Y.; and Liu, J. 2023. Symbol-LLM: Towards Foundational Symbol-centric Interface For Large Language Models. *arXiv preprint arXiv:2311.09278*.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; and Huang, F. 2024. Qwen2 Technical Report. *arXiv preprint arXiv:2407.10671*.
- Zhang, L.; Hu, A.; Xu, H.; Yan, M.; Xu, Y.; Jin, Q.; Zhang, J.; and Huang, F. 2024. TinyChart: Efficient Chart Understanding with Visual Token Merging and Program-of-Thoughts Learning. *arXiv preprint arXiv:2404.16635*.