

ViP-LLaVA: Making Large Multimodal Models Understand Arbitrary Visual Prompts

Mu Cai¹ Haotian Liu¹ Dennis Park² Siva Karthik Mustikovela²
 Gregory P. Meyer² Yuning Chai² Yong Jae Lee^{1,2}
¹University of Wisconsin–Madison ²Cruise LLC
<https://vip-llava.github.io>

Abstract

While existing large vision-language multimodal models focus on whole image understanding, there is a prominent gap in achieving region-specific comprehension. Current approaches that use textual coordinates or spatial encodings often fail to provide a user-friendly interface for visual prompting. To address this challenge, we introduce a novel multimodal model capable of decoding arbitrary (free-form) visual prompts. This allows users to intuitively mark images and interact with the model using natural cues like a “red bounding box” or “pointed arrow”. Our simple design directly overlays visual markers onto the RGB image, eliminating the need for complex region encodings, yet achieves state-of-the-art performance on region-understanding tasks like Visual7W, PointQA, and Visual Commonsense Reasoning benchmark. Furthermore, we present ViP-Bench, a comprehensive benchmark to assess the capability of models in understanding visual prompts across multiple dimensions, enabling future research in this domain. Code, data, and model are publicly available.

1. Introduction

Large language models (LLMs) like ChatGPT [32], GPT4 [33], and Bard [12] have recently gained significant attention for their strong reasoning and generalization capabilities, and their ability to chat in a human-like manner. In particular, models such as GPT-4V(ision) [31], which incorporate visual information, have demonstrated human-level perception and reasoning capabilities [50]. This has spurred the development of similar open-source models that aim to replicate or even surpass the proprietary models’ performance.

Despite their capabilities, current models, including seminal ones like LLaVA [23, 24] and MiniGPT-4 [56], focus predominantly on whole-image understanding; in other words, they lack the capability to process *region-specific* information in complex scenes. This limitation becomes par-

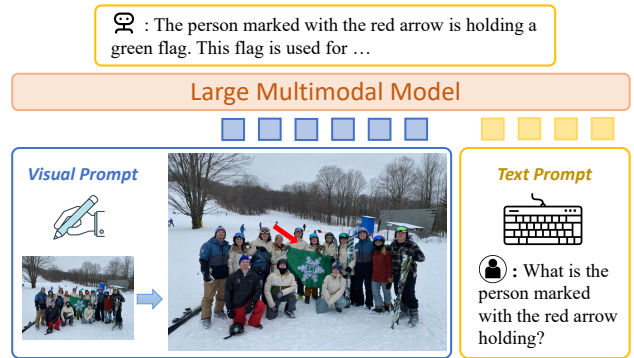


Figure 1. **Main Idea of ViP-LLaVA.** We directly overlay diverse visual prompts (e.g., arrows, boxes, circles, scribbles) onto the original image, and then feed the corresponding visual features along with text embeddings into the large multimodal model for conversational assistance. Here we show an example using a red arrow.

ticularly apparent when attempting to describe specific objects within an image using only language prompts, which can be difficult when there is ambiguity (e.g., when there are multiple people in the image, and the question relates to a specific person), as shown in Figure 1.

To address this issue, recent work explores spatial references in multimodal models. Existing efforts have primarily focused on using textual representations of coordinates [4, 5, 9, 53], learned positional embeddings [34, 52, 55], or ROI features [37, 52]. However, they often lack user-friendliness, as they are limited to fixed-format visual references like bounding boxes and the spatial coordinates of a mask contour. Most of these approaches, including those by Zhang et al. [52] and Chen et al. [5], only employ bounding box inputs for visual referrals. While effective in structured scenarios, this method proves less versatile in natural, user-driven interactions where the visual prompts may not conform to clean geometric shapes.

In this paper, we propose a simple yet highly effective solution to this problem: a large multimodal model that can

process *arbitrary visual prompts*. This allows a user to intuitively mark up images and interact using natural cues such as a “red bounding box” or “pointed arrow”. Our model recognizes these visual prompts, offering a user-friendly way to integrate visual references into the language dialogue. Based on our own observation and prior work [38], which shows that CLIP can understand visual markers, we *directly inject the visual prompts into the original image space* without any additional region-specific model designs. Although our approach is deceptively simple, it yields an unexpected benefit: our model sets new state-of-the-art performances on tasks demanding precise region-specific perception and complex reasoning. It surpasses the capabilities of existing related models with specialized region encoding techniques, as evidenced by our superior performance on region reasoning tasks on Visual7W [57] and PointQA [29].

To further support research in this area, we introduce *ViP-Bench*, a benchmark for evaluating multimodal models’ region understanding capabilities with arbitrary visual prompts. By collecting a diverse set of 303 images and questions, we provide a comprehensive assessment of visual understanding capabilities across six aspects at the region level: recognition, OCR, knowledge, math, object relationship reasoning, and language generation. We believe that ViP-Bench will provide a solid foundation for future research into multimodal models with arbitrary visual prompts.

In summary, our main contributions are:

- We introduce a novel multimodal model for intuitive interaction with images using natural language and arbitrary visual prompts, enhancing user accessibility and model flexibility.
- We develop a visual referral approach that overlays visual prompts directly onto images, simplifying the model’s architecture without compromising performance.
- Our model, ViP-LLaVA, achieves state-of-the-art results on region understanding tasks on established benchmarks, surpassing specialized region encoding models.
- We introduce ViP-Bench, a benchmark for evaluating visual prompt interpretation, setting a foundational platform for future research.

2. Related Work

Advancements in Large Multimodal Models. Large language models like ChatGPT [32], GPT4 [33], and LLaMA [41] have shown impressive reasoning and generalization capabilities. The landscape of LLMs has been markedly transformed by the recent introduction of models that integrate visual information, such as GPT-4V(ision) [31]. Building upon open-source LLMs [41, 43], a vast number of multimodal vision-language models have made significant strides, spearheaded by LLaVA [23, 24] and MiniGPT-4 [56], which combine LLaMA’s [41] lan-

guage prowess with a CLIP [36] based image encoder. While these models excel at whole-image understanding, a key challenge has been region-specific comprehension within complex visual scenes. This has led to the exploration of spatial referrals in multimodal contexts. Existing models utilize textual coordinate representations [4, 5, 9, 53], learned positional embeddings [34, 52, 55], or Region of Interest (ROI) features [52] to anchor language to specific image regions. However, they often employ rigid visual referral formats that are not as intuitive for users.

Visual Prompting as a User-Friendly Solution. Our focus is on making the interaction with multimodal models more natural and intuitive. Traditional models have employed regular shapes for visual prompting, but our research is motivated by the need for a system that can interpret a wider range of visual prompts. For example, in visual perception, interactive segmentation methods have been proposed that can take in points or scribbles [17, 58]. Drawing inspiration from recent findings that show GPT-4V’s ability to understand a variety of markers [46], we advocate for a model that can handle arbitrary visual cues, such as scribbles and arrows. In our model, ViP-LLaVA, we overlay these visual prompts directly onto the image canvas. This is accomplished by fine-tuning on a dataset specifically designed for arbitrary visual prompt instructions.

Evaluating LMM’s Region Understanding Capabilities.

Existing works [5, 34, 47, 52] evaluates the model’s region understanding capabilities on regional multichoice [29, 51, 57] or captioning [18, 49] tasks with metrics such as accuracy, recall, and CIDEr [42]. However, these metrics fall short when it comes to evaluating visual dialogue for large multimodal models in an open-world setting. To evaluate LMM’s capability in engaging in visual conversations for *image-level* understanding, two families of evaluation are proposed: multiple-choice [25] or using GPT4 as a judge for free-form answers [24, 50]. However, a gap still exists in the evaluation of LMM’s capabilities for comprehending arbitrary visual prompts. To address this, we introduce ViP-Bench, a comprehensive benchmark tailored to evaluate how well the LMMs can interpret various visual prompts across multiple dimensions, including recognition, OCR, knowledge, math, relationship reasoning, and language generation.

3. Approach

Our research hinges on the premise that a large multimodal model should not only perceive the visual content of an image but also interpret arbitrary visual markers as part of the user interaction. In this section, we describe our approach that achieves this goal, highlighting the pivotal role of CLIP in understanding visual markers and the construction of a

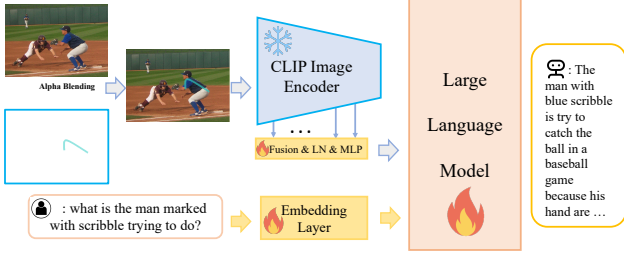


Figure 2. **Model Architecture.** After alpha blending the visual prompts onto the original image, we feed the resulting image into the visual encoder to obtain multi-level visual features. Those features are concatenated and fed into the LayerNorm and MLP layers to form the visual tokens. Then visual tokens and text instruction tokens are fed into the large language model to produce the language response in an auto-regressive manner. The frozen and trainable modules during instruction tuning are annotated.

new instruction tuning dataset tailored to train ViP-LLaVA to understand arbitrary visual prompts.

3.1. Visual Prompt Embedding via CLIP

In contrast to prior work on region understanding [34, 52] which constructs a new module to process visual prompts, we leverage CLIP’s [36] existing capabilities to encode both the image and superimposed visual markers. Specifically, CLIP’s proficiency in aligning visual and textual data makes it an ideal candidate for this task, as recent studies [38] suggest that it inherently pays attention to marked regions including circles, rectangles, *etc.* As shown in our experiments, we further demonstrate that CLIP can focus the model’s attention on a wider variety of visual prompts such as arrows and arbitrary scribbles. To utilize this functionality, we composite the visual prompts \mathbf{P}_v onto the original image \mathbf{X}_v through alpha blending, creating a merged representation that highlights the areas of interest:

$$\hat{\mathbf{X}}_v = \alpha \cdot \mathbf{P}_v + (1 - \alpha) \cdot \mathbf{X}_v, \quad (1)$$

where $\alpha \in [0, 1]$ denotes the transparency level of the visual prompt, \mathbf{X}_v is the image, and \mathbf{P}_v is the image with the visual prompt. Note that we only perform alpha blending for pixels underlying the visual prompt. The composite image $\hat{\mathbf{X}}_v$ is then fed into the multimodal model.

To effectively recognize the visual prompts, we balance low-level and high-level visual features in ViP-LLaVA.

To address the tendency of CLIP’s deeper features to overlook low-level details [54], we selectively extract features from multiple CLIP layers. Specifically, we use one early layer (6-th) to encode detailed geometric shapes and four deeper layers (15, 18, 21, 24-th) to capture broader semantic information. These multi-level features are then concatenated, normalized using LayerNorm [2] for training stability, and finally passed through an MLP layer. This

process ensures ViP-LLaVA effectively integrates diverse visual cues, a strategy validated through our ablation studies detailed in Sec. 5.4.

Our design’s simplicity of directly overlaying visual prompts offers several advantages. It reduces model complexity by bypassing additional processing modules and aligns closely with natural human interactions, as users often employ diverse and spontaneous visual markers. This flexibility allows ViP-LLaVA to interpret a wide range of user-generated visual cues, enhancing its applicability in real-world scenarios.

To train ViP-LLaVA, we perform autoregressive language modeling; *i.e.*, we maximize the likelihood of generating the tokens of the ground-truth answer \mathbf{X}_a :

$$P(\mathbf{X}_a | \hat{\mathbf{X}}_v, \mathbf{X}_{\text{instruct}}) = \prod_{i=1}^L P_{\theta}(x_i | \hat{\mathbf{X}}_v, \mathbf{X}_{\text{instruct}}, \mathbf{X}_{a, < i}) \quad (2)$$

where θ represents the trainable parameters, $\mathbf{X}_{\text{instruct}}$ is the text instruction, L is the sequence length of the answer \mathbf{X}_a , and $\mathbf{X}_{a, < i}$ denotes all the answer tokens before the current prediction token x_i , where i denotes the steps during text token generation. Here we omit system messages from the equation for clarity, even though they are part of the conditioning. Figure 2 shows our model architecture.

This training objective enables the model to generate contextually accurate responses by comprehending the visual content, language instruction, and the overlaid prompts. It fosters the model’s ability to interpret visual markers in unison with the image, thereby enhancing its proficiency in addressing complex, region-specific language inquiries. This capability is crucial for tasks requiring nuanced understanding of both the visual elements and user intentions conveyed through arbitrary visual prompts.

3.2. Visual Prompting Design

To train the model to recognize and interpret arbitrary visual prompts, we develop a new visual prompt instruction tuning dataset, as there are no prior datasets with arbitrary visual prompts and instruction-output text pairs that we can use.

Our dataset comprises a diverse collection of 520k image-text pairs marked with visual prompts, sourced from publicly available datasets, including (1) single region reasoning data: 80k referring comprehension and generation data from RefCOCOg [49], and 37k object counting data from PointQA-LookTwice [29], (2) two-region reasoning data: 80k triplet relationship data from Visual Genome [18], (3) multi-region reasoning data: 30k grounded image captioning data from Flickr 30k Entities [35], 213K data from Visual Commonsense Reasoning dataset [51], and 82k data from Visual7W [57]. Note that all those data are collected from the training split of the aforementioned datasets.

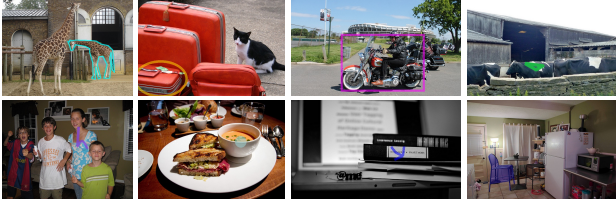


Figure 3. **Visualization of Visual Prompt Types.** From top-left to bottom-right: mask contour, ellipse, bounding box, triangle, scribble, point, arrow, and mask. Note that the prompts not only have diverse shapes, but they also have diverse colors, transparency values, widths, scales, and directions.

We automatically annotate each image with various visual prompts. For the data that only comes with bounding box annotations, we sample the visual prompts from three possible categories: rectangle, ellipse, and arrow. Here we make sure that the head of the arrow lies within $[(-\frac{W}{2}, -\frac{H}{2}), (\frac{W}{2}, \frac{H}{2})]$ space, where W, H are the width and height of the image, respectively. For ellipse, the lengths along the semi-major and semi-minor axes are inherited from the bounding box size, where we enlarge the ellipse with a ratio between $[1, 1.5]$. On the other hand, for regions that come with ground truth pixel-level mask annotations, we annotate each region with visual prompts sampled from the following 8 possibilities: rectangle, ellipse, point, triangle, mask, mask contour, arrow, and scribble created using Bézier curves; see Figure 3. We make sure that the head of the arrow, entire point, triangle, and scribble lies within the provided mask. These annotations simulate natural human interactions with images, where users often use spontaneous markers to highlight areas of interest.

For scribbles, we simulate human-like drawings using Bézier curves [8]. This process begins by randomly selecting three points within the object mask, which serve as the anchors for the quadratic Bézier curve. The generated Bézier curve is then composited onto the image using the previously mentioned alpha blending technique to produce a merged image with the scribble serving as a visual prompt.

Humans naturally use various markers to highlight objects within their environment. For instance, in educational settings, teachers often use arrows or underlining to draw students’ attention to specific parts of an image or text. Similarly, in everyday communication, people might circle items in a photograph to point out something of interest or use scribbles to obscure sensitive information before sharing. Through our design, we create a visual instruction following dataset that mirrors the way humans visually interact with objects, thus fostering a more intuitive and natural interaction with the model.

3.3. Optional Region-level Instruction Tuning Data

Our training data comes from two sources: (i) region-level visual prompting data described in Section 3.2, and (ii)

image-level data devoid of visual prompts, sourced from LLaVA-1.5 [23]. This strategy enables ViP-LLaVA to engage in human-like conversations, primarily due to the image-level LLaVA instruction data from Liu *et al.* [24]. Optionally, to further enhance ViP-LLaVA’s capability in multimodal conversations at the region-level, we design region-specific instruction data with the help of GPT-4V.

Prior approaches like Shikra [5] attempted to generate region-level instruction data using text-only models like GPT4. However, this method is inherently limiting, particularly in object-level tasks where the model, lacking visual context, cannot accurately reference multiple objects of the same class within a single scene. To overcome this, we develop an instruction data curation method using GPT-4V. Unlike text-only models, GPT-4V can interpret visual prompts displayed in images [46]. Our method involves feeding two images into GPT-4V: the original image and a modified version with annotated visual prompts. Alongside these images, we provide the model with the ground-truth (text) annotation in the original dataset and system messages. This process is used to curate `<visual prompt, text prompt, text output>` triplets for the images in our dataset described in Section 3.2.

We introduce specific textual representations such as `<within red mask>` and `(<within red box>, <within blue box>)` to guide GPT-4V in recognizing the visual prompts in both single-region and multi-region settings. During training, we replace these phrases with the set of eight possible visual prompts described in Section 3.2, significantly enhancing the dataset’s versatility. In total, we curate 13k high-quality region-level instruction data points, comprised of 7k single-region and 6k multi-region instances. In the supplementary, we provide specific details of the system messages, input text prompts, and generated text outputs.

Although ViP-LLaVA works well even without this enriched data for standard visual reasoning benchmarks, we find that it helps to further improve the model’s ability to have human-like conversations in open-world settings.

4. ViP-Bench for Evaluation

In order to rigorously evaluate the capabilities of multimodal models in interpreting and responding to visual reasoning queries, we introduce ViP-Bench, a benchmarking suite for evaluating multimodal region-understanding capabilities under various visual prompts. ViP-Bench consists of 303 unique image-question pairs, where images are collected from MM-Vet [50], MMBench [25], and Visual Genome [18]. Each pair consists of an image coupled with a diverse visual reasoning question designed to test a model’s understanding and interpretation capabilities. We reuse the questions in MM-Vet [50] and MMBench [25] (but make minor adjustments so that they take into account the region-

Method	Generalist?	Accuracy (%)
LSTM-Att [57]	×	56.10
CMNs [15]	×	72.53
12in1 [27]	×	83.35
GPT4ROI-7B [52]	×	81.83
GPT4ROI-13B [52]	×	84.82
Shikra-13B [5]	✓	85.33
ViP-LLaVA-Base-7B	✓	86.04
ViP-LLaVA-Base-13B	✓	87.54
ViP-LLaVA-7B	✓	86.60
ViP-LLaVA-13B	✓	87.91

Table 1. Comparison of methods in terms of generality and accuracy on Visual7W [57] test set.

specific visual prompts), while in Visual Genome, we design the questions and answers by ourselves. We use bounding boxes and masks produced by the Segment Anything Model (SAM) [17] to annotate the location of the objects.

Key to the design of ViP-Bench is its comprehensive coverage of six crucial aspects of visual understanding at the region level: recognition, OCR (Optical Character Recognition), knowledge, math, object relationship reasoning, and language generation. This range ensures a holistic assessment of a model’s performance in various facets of region-level visual reasoning.

ViP-Bench employs a similar grading mechanism as MM-Vet [50]. We employ the GPT-4 text model, a state-of-the-art language model, to evaluate the responses of multimodal models. Specifically, we feed the response from the multimodal model, the human annotated answer, and several in-context scoring examples to GPT-4. The responses are scored by GPT-4 on a scale from 0 to 10, offering a quantitative measure of the multimodal model’s proficiency in understanding and interpreting visual data. This grading system provides a standardized framework for comparing the performance of different models.

ViP-Bench is meticulously annotated by humans. This process involved seven rounds of validation to ensure the accuracy and relevance of the object boxes/masks, questions, and answers. Such rigorous annotation guarantees the reliability of the benchmark as a tool for model evaluation. An illustrative example in Table 6 showcases a scenario where a leading model like GPT-4V misinterprets object localization under ViP-Bench, highlighting the challenges in current multimodal understanding. We present additional visualizations and statistics of ViP-Bench in the supp.

Through ViP-Bench, we provide a valuable tool for the research community, aiding in the development and refinement of multimodal models. By offering a comprehensive and challenging testbed, we believe ViP-Bench can set the stage for future advancements in the field of visual reasoning and multimodal interaction.

Method	Generalist?	Accuracy (%)
Point and ask [29]	×	60.20
LLaVA-1.5-7B [23]	✓	56.19 [†]
LLaVA-1.5-13B [23]	✓	57.93 [†]
Shikra-13B [5]	✓	70.30
ViP-LLaVA-Base-7B	✓	70.86
ViP-LLaVA-Base-13B	✓	72.15
ViP-LLaVA-7B	✓	71.31
ViP-LLaVA-13B	✓	71.77

Table 2. Comparison of methods in terms of generality and accuracy on PointQA-LookTwice [29] test set. [†]zero-shot eval.

5. Experiments

In this section, we compare ViP-LLaVA to state-of-the-art multimodal models, including those that explicitly design region-specific modules, perform in-depth analysis to assess ViP-LLaVA’s capabilities, and perform ablation studies.

5.1. Training Setup

Model. For the visual model, we choose CLIP-336px [36] to preserve more information from the raw pixel space. We use Vicuna v1.5 [43] as the language encoder. For the multimodal connector, a 2-layer MLP is utilized. Ablations on more LLM backbones are shown in Supp. A.5.

Training and data. During the initial stage of training, we employ 558k BLIP [6, 24] captioned image-text pairs to pretrain the multimodal connector. The second stage utilizes LLaVA v1.5 [23] instruction data alongside our region-level visual prompting dataset from Section 3.2. Both stages train the model for 1 epoch, with an overall training time of around 20/40 hours for the 7B/13B model using 8 NVIDIA A100 GPUs. Finally, we mix the 13k GPT-4V instruction data with 13k sampled data from stage 2 to get 26k stage 3 training data, and then fine-tune our stage-2 model (referred to as ViP-LLaVA-Base) for one epoch to get our model ViP-LLaVA, which requires approximately 0.5 hours for the 7B model and 1 hour for the 13B model on 8 NVIDIA A100 GPUs.

Visual prompts. ViP-LLaVA uses 8 visual prompts: rectangles, ellipses, points, scribbles, triangles, masks, mask contours, and arrows. Their attributes, such as color, thickness, and alpha value for alpha blending (in [0.5, 1]) are randomized. The arrow’s direction and length are randomized, with the endpoint remaining within the mask. For referencing specific regions, we replace the <region> text with the color and shape description, such as *red scribble*. The visual prompt type and associated attributes for each region are randomly assigned during training.

Model	$Q \rightarrow A$ (%)	$QA \rightarrow R$ (%)	$Q \rightarrow AR$ (%)
ViLBER [26]	72.4	74.5	54.0
Unicoder-VL [20]	72.6	74.5	54.5
VLBERT-L [40]	75.5	77.9	58.9
ERNIE-ViL-L [48]	78.52	83.37	65.81
VILLA-L [11]	78.45	82.57	65.18
GPT4RoI-7B [52]	87.4	89.6	78.6
ViP-LLaVA-Base-7B	87.66	89.80	78.93

Table 3. Validation Accuracy on VCR [51] dataset.

5.2. Evaluation on Region Reasoning Benchmarks

We first quantitatively evaluate ViP-LLaVA on three region reasoning benchmarks.

Visual7W. The Visual7W dataset [57] tests models’ spatial perception by requiring them to match text descriptions with the correct bounding boxes from a set of choices. We differentiate between ‘generalist’ models, which are not specifically trained on the target dataset, and ‘specialist’ models, which are. For a fair comparison, we use image overlays as visual prompts for the LLaVA model and textual coordinates for Shikra’s text prompts. The results in Table 1 shows ViP-LLaVA-7B outperforming recent state-of-the-art methods, including GPT4RoI [52] and Shikra [5], despite having fewer parameters, and ViP-LLaVA-13B producing even higher gains. ViP-LLaVA overlays bounding boxes directly onto the image, creating an immediate link between the image and spatial locations. This contrasts with other methods that rely on external embeddings from either textual or newly learned embedding spaces to reference specific regions, proving less effective in this context.

PointQA-LookTwice. PointQA [29] presents a dataset where queries are based on either a specific point or a bounding box within an image. We evaluate ViP-LLaVA under the broad-question scenario using the bounding box type, typified by the prompt *How many of these are there?* This requires the model to first correctly identify the object within the given region and subsequently enumerate instances of the same category across the image—essentially a test of object recognition followed by class-specific counting. In line with our methodology for Visual7W, we use the image overlaid with the bounding box for LLaVA, while for Shikra, we incorporate the bounding box coordinates into the text prompt. Table 2 shows ViP-LLaVA’s superior performance on this intricate task, surpassing other multimodal contenders. Our method of overlaying visual prompts ensures the object remains unobscured, effectively combining the original image pixels with visual cues to enhance object recognition and counting accuracy.

Visual Commonsense Reasoning. The Visual Commonsense Reasoning (VCR) dataset [51] is a challenging bench-

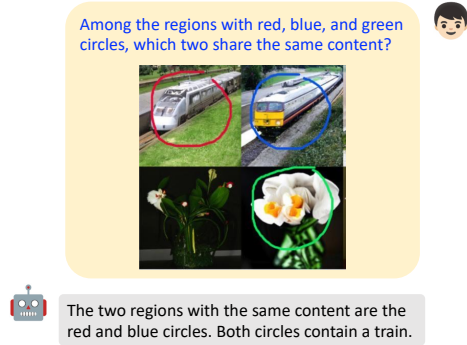


Figure 4. ViP-LLaVA model is able to infer correspondences between multiple objects in the image.

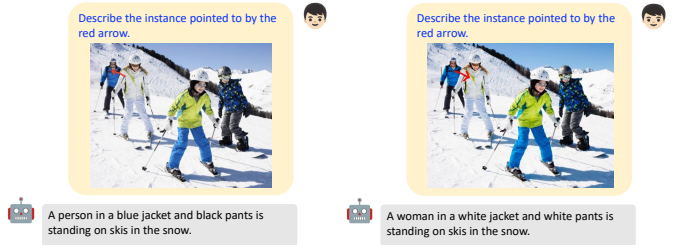


Figure 5. ViP-LLaVA is able to understand the direction of arrows.

mark designed to evaluate a model’s capabilities in high-level cognition and commonsense reasoning in the context of visual information. The dataset presents multiple-choice questions that require an understanding of the scene depicted in an image. Each question (Q) is paired with four potential answers (A), where the model must not only select the correct answer but also provide a rationale (R) that justifies its choice, demonstrating the model’s ability to comprehend and rationalize visual elements within a given context.

We finetune ViP-LLaVA-Base-7B on VCR, similar to the protocol in GPT4RoI [52]. As shown in Table 3, our approach exhibits state-of-the-art performance on the validation set, illustrating its proficiency in visual commonsense reasoning. This success highlights our approach’s dual strengths: adeptness in perception tasks and effectiveness in multi-region reasoning. By integrating visual prompts directly into the image, our model more effectively associates spatial locations with semantic understanding, facilitating a better interaction between spatial and semantic reasoning.

5.3. In-depth Analysis

Region reasoning with arbitrary human drawings at test time. ViP-LLaVA, when presented with arbitrarily drawn enclosed regions or arrows by a user, can accurately describe, shown in Figure 4 and 6.

Multi-region understanding capabilities. ViP-LLaVA demonstrates robust multi-region understanding, able to

Model	Visual7W	PointQA	VCR ($Q \rightarrow A$) (%)
Single Layer	82.44	70.23	84.64
Multi-Layer	86.04	70.86	87.66

Table 4. Single layer vs. multi-layer CLIP image encodings.

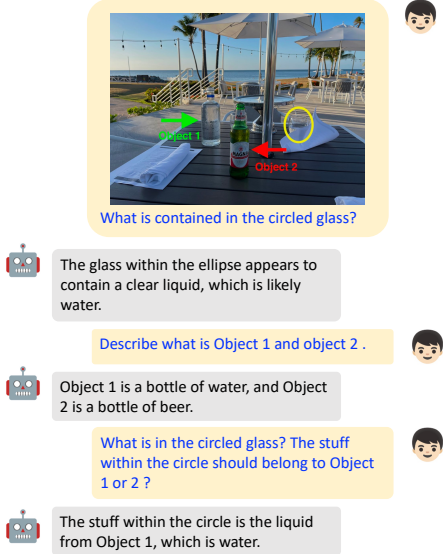


Figure 6. ViP-LLaVA is able to understand the text markers in the image, and reason about them.

dissect complex visual scenes and infer relationships between various elements. As shown in Figure 4, ViP-LLaVA is able to infer correspondences between multiple objects in the image, and make the correct reasoning that the red and blue circles both include the train.

Arrow direction understanding. ViP-LLaVA is able to understand arrows. Here we conduct an ablation study of the arrow direction. Given two arrows that have the same body yet different heads, as shown in Figure 5, ViP-LLaVA is able to understand the direction of the arrows, making correction descriptions about the respective regions.

Generalization to other attributes. ViP-LLaVA also generalizes to untrained attributes, like varying visual prompt thickness or location, showcasing its adaptability beyond what was seen during training. See the supplementary material for examples of different thicknesses.

Figure 6 further shows that ViP-LLaVA is able to conduct OCR first, and then make correspondences between different regions to make a correct prediction about the content of each part.

5.4. Ablation Studies

Impact of overlaying visual prompts on visual information. To assess whether overlaying visual prompts on images obscures visual information, we conduct a comparison by inputting visual tokens from both the original and over-

laid images into ViP-LLaVA-Base-7B. Using the VCR dataset, we evaluate the accuracy of the QA task with and without the additional visual tokens from the original image. Results on the VCR validation split shows an accuracy of 81.63% with the original image and overlaid image tokens, compared to 82.47% with the overlaid image tokens only. The similar accuracies suggest that the overlaid prompts do not detract from the visual information processed by our model.

Influence of CLIP multi-layer features. We next explore the impact of using multi-layer visual features from CLIP as opposed to single-layer features, specifically focusing on the second-last layer as implemented in LLaVA [23, 24]. Our ablation study in Table 4 reveals a marked improvement in performance, particularly in scenarios involving multiple visual prompts, as in the Visual7W and VCR datasets. This indicates that leveraging multi-layer visual features significantly enhances the model’s ability to localize and recognize visual prompts within images.

6. ViP-Bench Evaluation Results

Finally, we evaluate on ViP-Bench using a set of image-level and region-level LMMs, including InstructBLIP [6], GPT-4V [31], LLaVA v1.5 [24], Qwen-VL [3], Shikra [5], GPT4ROI [52] and Kosmos-2 [34]. For open-source models, we evaluate with greedy decoding (temperature=0). As shown in Table 5, we first see that the performance of all models, including GPT-4V, is far from perfect, demonstrating the difficulty of ViP-Bench. An illustrative case in Table 6 depicts a scenario where GPT-4V and LLaVA incorrectly predict object localization. Overall, ViP-LLaVA outperforms other models, except GPT-4V, demonstrating greater adaptability to various visual perception and reasoning tasks. By training on images overlaid with visual prompts, ViP-LLaVA becomes adept at understanding arbitrary visual cues and mimicks the natural human method of referring to objects in images. This enables it not only to better identify and interpret visual prompts but also to integrate these prompts into its reasoning process, enhancing its overall comprehension and response accuracy.

Visual prompting is superior to other representations.

In zero-shot evaluation, when visual prompts are represented as a simple list of four textual numerical values, models like Qwen-VL and LLaVA underperform compared to ViP-LLaVA. This underscores the effectiveness of visual prompts over basic textual representations.

Language tasks: A challenge for current LMMs. The ViP-Bench results reveal that, compared to GPT-4V, open-source LMMs show a significant gap in OCR, math, and language generation tasks, while they perform decently in recognition, knowledge, and object relationship reasoning.

Model	Format	Synthesized visual prompts (tight bounding box)							Visual prompts from human (e.g. arrow, circle)						
		Rec	OCR	Know	Math	Rel	Lang	All	Rec	OCR	Know	Math	Rel	Lang	All
GPT-4V-turbo-detail:high [33]	VP	58.1	69.8	59.5	71.0	61.4	51.9	60.7	56.9	69.7	63.7	80.6	61.1	45.6	59.9
GPT-4V-turbo-detail:low [33]	VP	53.2	50.3	55.6	67.7	57.5	57.5	52.8	51.7	50.3	59.3	60.3	55.0	43.8	51.4
InstructBLIP-7B [6]	VP	36.9	16.3	34.2	22.3	26.8	7.5	31.7	38.9	17	35.4	9.7	29.3	17.5	33.3
Shikra 7B [5]	Coor	40.2	10.0	28.0	3.5	18.9	20.6	33.7	-	-	-	-	-	-	-
GPT4ROI 7B [52]	ROI	35.6	16.7	29.7	9.7	32.5	13.8	35.1	-	-	-	-	-	-	-
Kosmos-2 [34]	Dis	29.5	14.2	18.5	9.7	7.5	21.9	26.9	-	-	-	-	-	-	-
LLaVA-1.5-7B [24]	Coor	52.7	20.7	44.7	14.5	44.6	30.6	44.8	-	-	-	-	-	-	-
LLaVA-1.5-7B [24]	VP	50.8	12.4	49.2	6.5	51.8	23.8	41.6	49.1	13	42.9	9.7	50	27.5	40.2
Qwen-VL-Chat [3]	Coor	52.6	22.0	40.0	12.9	47.1	26.9	45.3	-	-	-	-	-	-	-
Qwen-VL-Chat [3]	VP	43.0	30.4	40.2	9.7	25.7	28.7	39.2	48.7	22.1	41.2	6.5	48.2	25	41.7
ViP-LLaVA-Base-7B	VP	54.8	18.8	52.9	9.7	53.9	42.5	45.5	55.3	17.6	45.9	8.1	44.6	33.1	46.8
ViP-LLaVA-7B	VP	56.7	19.4	49.7	10.0	50.4	33.8	48.4	56.7	21.2	47.1	12.3	50.4	36.2	48.3
InstructBLIP-13B [6]	VP	42.5	12.2	37.5	3.2	33.2	12.5	35.8	41.7	13.6	35.9	3.2	27.9	18.8	35.2
LLaVA-1.5-13B [24]	Coor	53.2	26.1	45.9	9.7	52.5	31.9	47.1	-	-	-	-	-	-	-
LLaVA-1.5-13B [24]	VP	48.1	21.8	40.0	6.1	45	28.1	41.8	48.8	21.3	47.3	15.8	44.6	31.2	42.9
ViP-LLaVA-Base-13B	VP	54.4	27.8	51.2	16.1	51.1	46.9	48.2	55.3	21.1	46.8	9.7	45.4	38.8	47.0
ViP-LLaVA-13B	VP	56.3	24.6	53.4	15.5	50.0	53.8	48.3	55.4	26.9	49.3	15.5	48.6	41.9	48.2

Table 5. **ViP-Bench Evaluation Results.** This table presents the performance of various models under ViP-Bench, utilizing different visual prompt formats. The evaluation includes both synthesized and human-drawn prompts, providing insights into the models’ maximum potential and real-world applicability, respectively. Formats include VP (visual prompts), Coor (coordinates as visual prompts), Dis (discrete positional tokens for vocabulary expansion), and ROI (CLIP region of interest features with positional embedding). The assessed dimensions are Recognition (Rec), OCR, Knowledge (Know), Math, Relationship (Rel), and Language Generation (Lang).

This suggests that future VLM developments should prioritize enhancing language reasoning capabilities. For OCR, the results indicate a need for higher resolution inputs or a more robust backbone model, moving beyond the existing capabilities of models like CLIP.

Overfitting Concerns in Region-Level LMMs. Current region-level LMMs, including Shikra [5], GPT4ROI [52] and Kosmos-2 [34], tend to struggle with tasks involving mathematics, relationship reasoning, and language generation. This trend suggests a potential overfitting issue with these models to existing public region-level datasets, which predominantly feature brief descriptions.

7. Conclusion

In summary, ViP-LLaVA shows that visual prompts are promising for region-specific image understanding. By integrating arbitrary visual prompts, we bridge the gap between user-friendly interfaces and the precision required for region comprehension. ViP-LLaVA’s intuitive design leverages natural linguistic interactions coupled with visual markers, simplifying the process of image annotation while enhancing the clarity of visual references. Our state-of-the-art performance on established benchmarks including Visual7W, PointQA, and VCR, underlines the efficacy of ViP-LLaVA. Notably, the introduction of ViP-Bench as a comprehensive evaluative platform sets a new standard for assessing multimodal models’ region reasoning abilities. ViP-LLaVA establishes a foundation for further exploration

Visual input example for Spatial Reasoning:



User	Between Object 1: the object within the red mask contour, Object 2: the object within the blue mask contour, and Object 3: the object within the green mask contour, which one has something on top of it?
Ground Truth	Object 2
GPT-4V	Object 3, the object within the green mask contour, has something on top of it. It appears to be a coffee pot placed on top of what could be a coffee machine.
LLaVA-1.5-13B	Object 3, which is the coffee maker, has something on top of it.
ViP-LLaVA-13B	Object 2: the object within the blue mask contour has something on top of it.

Table 6. An example in ViP-Bench where GPT-4V makes a wrong prediction. The correct answer should be Object 2.

in the field of intelligent visual systems. We believe that ViP-LLaVA can motivate how visual and linguistic modalities are integrated, enabling more sophisticated and nuanced human-machine interactions.

Acknowledgements. This work was supported in part by NSF CAREER IIS2150012, and Institute of Information &

communications Technology Planning & Evaluation(IITP) grants funded by the Korea government(MSIT) (No. 2022-0-00871, Development of AI Autonomy and Knowledge Enhancement for AI Agent Collaboration). (No. RS2022-00187238, Development of Large Korean Language Model Technology for Efficient Pre-training), and Microsoft Accelerate Foundation Models Research Program.

References

- [1] Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024. [4](#)
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [3](#)
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023. [7](#), [8](#)
- [4] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechu Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*, 2023. [1](#), [2](#)
- [5] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#), [3](#)
- [6] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. [5](#), [7](#), [8](#)
- [7] Michael Denkowski and Alon Lavie. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*, 2014. [5](#), [8](#)
- [8] Gerald Farin. *Curves and Surfaces for Computer-Aided Geometric Design: A Practical Guide*. Academic Press, 2014. [4](#)
- [9] Jon Ferraiolo, Fujisawa Jun, and Dean Jackson. *Scalable vector graphics (SVG) 1.0 specification*. iuniverse Bloomington, 2000. [1](#), [2](#)
- [10] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. [4](#)
- [11] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. *Advances in Neural Information Processing Systems*, 33:6616–6628, 2020. [6](#)
- [12] Google. Google bard. <https://bard.google.com/chat/>, 2023. [1](#)
- [13] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. [4](#)
- [14] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617, 2018. [4](#)
- [15] Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. Modeling relationships in referential expressions with compositional modular networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1115–1124, 2017. [5](#)
- [16] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 2019. [4](#)
- [17] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. [2](#), [5](#)
- [18] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017. [2](#), [3](#), [4](#)
- [19] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023. [4](#)
- [20] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11336–11344, 2020. [6](#)
- [21] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023. [4](#)
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [2](#)
- [23] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. [1](#), [2](#), [4](#), [5](#), [7](#), [6](#)
- [24] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv:2304.08485*, 2023. [1](#), [2](#), [4](#), [5](#), [7](#), [8](#)
- [25] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He,

- Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023. [2](#), [4](#)
- [26] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vlb- bert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019. [6](#)
- [27] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10437–10446, 2020. [5](#)
- [28] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 2022. [4](#)
- [29] Arjun Mani, Nobline Yoo, Will Hinthorn, and Olga Russakovsky. Point and ask: Incorporating pointing into visual question answering. *arXiv preprint arXiv:2011.13681*, 2020. [2](#), [3](#), [5](#), [6](#), [4](#)
- [30] Meta. Llama-3. <https://ai.meta.com/blog/meta-llama-3/>, 2024. [4](#)
- [31] OpenAI. Gpt-4v(ision) system card. https://cdn.openai.com/papers/GPTV_System_Card.pdf, 2023. [1](#), [2](#), [7](#)
- [32] OpenAI. Chatgpt. <https://openai.com/blog/chatgpt/>, 2023. [1](#), [2](#)
- [33] OpenAI. Gpt-4 technical report. 2023. [1](#), [2](#), [8](#)
- [34] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023. [1](#), [2](#), [3](#), [7](#), [8](#)
- [35] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. [3](#)
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#), [3](#), [5](#)
- [37] Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Erix Xing, Ming-Hsuan Yang, and Fahad S Khan. Glamm: Pixel grounding large multimodal model. *arXiv preprint arXiv:2311.03356*, 2023. [1](#), [8](#)
- [38] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. What does clip know about a red circle? visual prompt engineering for vlms. *ICCV*, 2023. [2](#), [3](#)
- [39] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. [4](#)
- [40] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019. [6](#)
- [41] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. [2](#)
- [42] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. [2](#), [5](#), [8](#)
- [43] Vicuna. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. <https://vicuna.lmsys.org/>, 2023. [2](#), [5](#), [4](#)
- [44] Jialian Wu, Jianfeng Wang, Zhengyuan Yang, Zhe Gan, Zicheng Liu, Junsong Yuan, and Lijuan Wang. Grit: A generative region-to-text transformer for object understanding. *arXiv preprint arXiv:2212.00280*, 2022. [8](#)
- [45] Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*, 2023. [1](#)
- [46] Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. The dawn of lmms: Preliminary explorations with gpt-4v(ision). *arXiv preprint arXiv:2309.17421*, 2023. [2](#), [4](#)
- [47] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023. [2](#)
- [48] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vil: Knowledge enhanced vision-language representations through scene graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3208–3216, 2021. [6](#)
- [49] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 69–85. Springer, 2016. [2](#), [3](#), [5](#)
- [50] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023. [1](#), [2](#), [4](#), [5](#)
- [51] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#), [3](#), [6](#)

- [52] Shilong Zhang, Peize Sun, Shoufa Chen, Min Xiao, Wenqi Shao, Wenwei Zhang, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv preprint arXiv:2307.03601*, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [53] Liang Zhao, En Yu, Zheng Ge, Jinrong Yang, Haoran Wei, Hongyu Zhou, Jianjian Sun, Yuang Peng, Runpei Dong, Chunrui Han, et al. Chatspot: Bootstrapping multimodal llms via precise referring instruction tuning. *arXiv preprint arXiv:2307.09474*, 2023. [1](#), [2](#)
- [54] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision (ECCV)*, 2022. [3](#)
- [55] Qiang Zhou, Chaohui Yu, Shaofeng Zhang, Sitong Wu, Zhibing Wang, and Fan Wang. Regionblip: A unified multimodal pre-training framework for holistic and regional comprehension, 2023. [1](#), [2](#)
- [56] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. [1](#), [2](#)
- [57] Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7W: Grounded Question Answering in Images. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. [2](#), [3](#), [5](#), [6](#), [4](#)
- [58] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *NeurIPS*, 2023. [2](#)

Supplementary Material

This supplementary document extends our main paper by providing additional results and in-depth analyses that were not included in the main manuscript due to space limitations. In Section A, we offer both qualitative and quantitative analyses on topics such as visual prompt generation, effect of the instruction data, arrow direction understanding, perform under each visual prompt, and impacts of different LLMs on the conventional vision-language model benchmarks, thus providing a comprehensive examination of our research. In Section B, we include more training details. Section C delves into further ablation studies, illuminating the design and efficacy of our approach. Additional experimental results, including a focus on region captioning, are discussed in Section D. The exploration of visual prompt augmentation’s potential is presented in Section E. Lastly, Section F comprises detailed statistics and visualizations from the ViP-Bench dataset.

A. In-Depth Analysis

A.1. Generalization to New Attributes

ViP-LLaVA, having been trained on eight types of visual prompts—namely mask contour, ellipse, bounding box, triangle, scribble, point, arrow, and mask—exhibits notable generalization capabilities. In the main paper, we show that ViP-LLaVA can understand human drawn visual prompts, as shown in Figure 7. As demonstrated in Figures 8 and Figure 5 of the main paper, ViP-LLaVA adeptly handles visual prompts with varying thicknesses and diverse markers, even though it was not explicitly trained on such variations. Furthermore, it effectively interprets text markers as visual prompts, a feature inspired by the Set-of-Mark [45].

Figures 9, 10, and 11 present qualitative examples. In Figure 9, ViP-LLaVA accurately localizes objects tagged with the digits “1”, “2”, and “3”, and generates precise descriptions for each. Figure 10 showcases the model’s ability to recognize digit markers and describe the color of vehicles accurately, despite the markers displaying counterfactual colors relative to the actual vehicle colors. Figure 11 illustrates the model’s competency in localizing a lemon within a scene densely populated with markers.

A.2. Effect of Optional GPT-4V Region-Level Instruction Data

As mentioned in Section 3.3 of the main paper, incorporating GPT-4V as an additional source of instruction data can enhance ViP-LLaVA’s performance. An example of the curation process is shown in Figure 12. For this purpose, we combine 13K data entries from the original stage 2 instruction dataset with an equal number of GPT-4V region-level

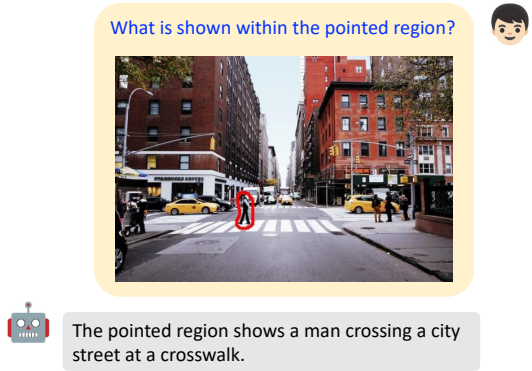


Figure 7. ViP-LLaVA is able to recognize a tiny region specified by the red contour.

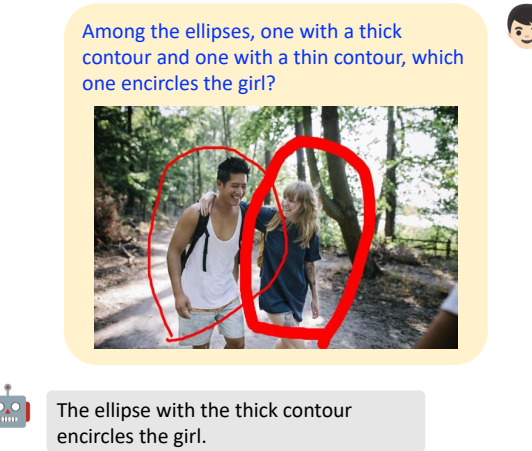


Figure 8. ViP-LLaVA is able to recognize visual prompts with different thickness not seen during training.

instruction data entries, forming a comprehensive 26K-entry stage 3 fine-tuning dataset. We fine-tune our stage-2 model for one epoch, which requires approximately 0.5 hours for the 7B model and 1 hour for the 13B model on 8 NVIDIA A100 GPUs. As shown in Table 8, the fine-tuned model, designated as ViP-LLaVA, demonstrates improvements across nearly all datasets for both the 7B and 13B models, underscoring the efficacy of the GPT-4V instruction data curation process. Notably, even without the GPT-4V instruction data, ViP-LLaVA outperforms contemporary methods on benchmarks such as Visual7W, PointQA-LookTwice, and ViP-Bench. The inclusion of GPT-4V instruction data further amplifies this performance advantage.

(a)

Solve the following equations:

$$1) 8x + 11 = 4x + 14$$

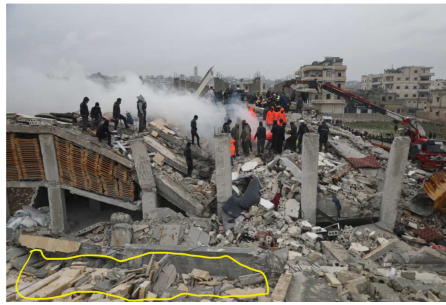
$$2) 7d - 4 = 11d - 9$$

Q: What is the value of the variable in the equation within the blue ellipse?

GT: 0.75 OR $x=0.75$ OR $3/4$

Required capabilities: OCR, MATH

(b)



Q: Are the debris within the yellow mask contour made of the same material?

GT: No

Required capabilities: Recognition, Knowledge

(c)



Q: Is the person pointed by red arrow going to be mad at the person pointed by blue arrow, for looking at the person pointed by green arrow? Answer it and give the rationale.

GT: Yes, the woman in red and the man appear to be a couple and the woman in red would not appreciate the man checking out other women. I think so because in a monogamous relationship, partners are expected to be faithful to one another.

Required capabilities: Recognition, Knowledge, Language Generation, Relationship

Table 7. Three samples requiring different capability integrations.

A.3. Understanding Arrow Direction

To rigorously evaluate ViP-LLaVA’s capacity for interpreting arrow directions, we next construct a challenging dataset of examples derived from the COCO validation set [22]. Specifically, we generate multiple scenarios

with arrows: each arrow originates from the center of one object’s bounding box and points towards the center of another, and vice versa. These visualizations are depicted in Figure 13. The typical prompt used is as follows: Determine whether object A (**category1**)

Model	Use GPT-4V Data?	Visual7W	PointQA-LookTwice	ViP-Bench@Box	ViP-Bench@Human
GPT4ROI-7B [52]	×	81.83 [†]	–	35.1	–
Shikra-7B [5]	×	–	–	33.7	–
ViP-LLaVA-Base-7B	×	86.04	70.86	45.5	46.8
ViP-LLaVA-7B	✓	86.09	71.31	48.4	48.3
GPT4ROI-13B [52]	×	84.82 [†]	–	–	–
Shikra-13B [5]	×	85.33	70.30	–	–
ViP-LLaVA-Base-13B	×	87.54	72.15	48.2	47.0
ViP-LLaVA-13B	✓	88.28	71.77	48.3	48.2

Table 8. Comparative performance analysis of the use of GPT-4V data in the 7B and 13B models. [†] Indicates that GPT4ROI specifically trained a specialist model on Visual7W.

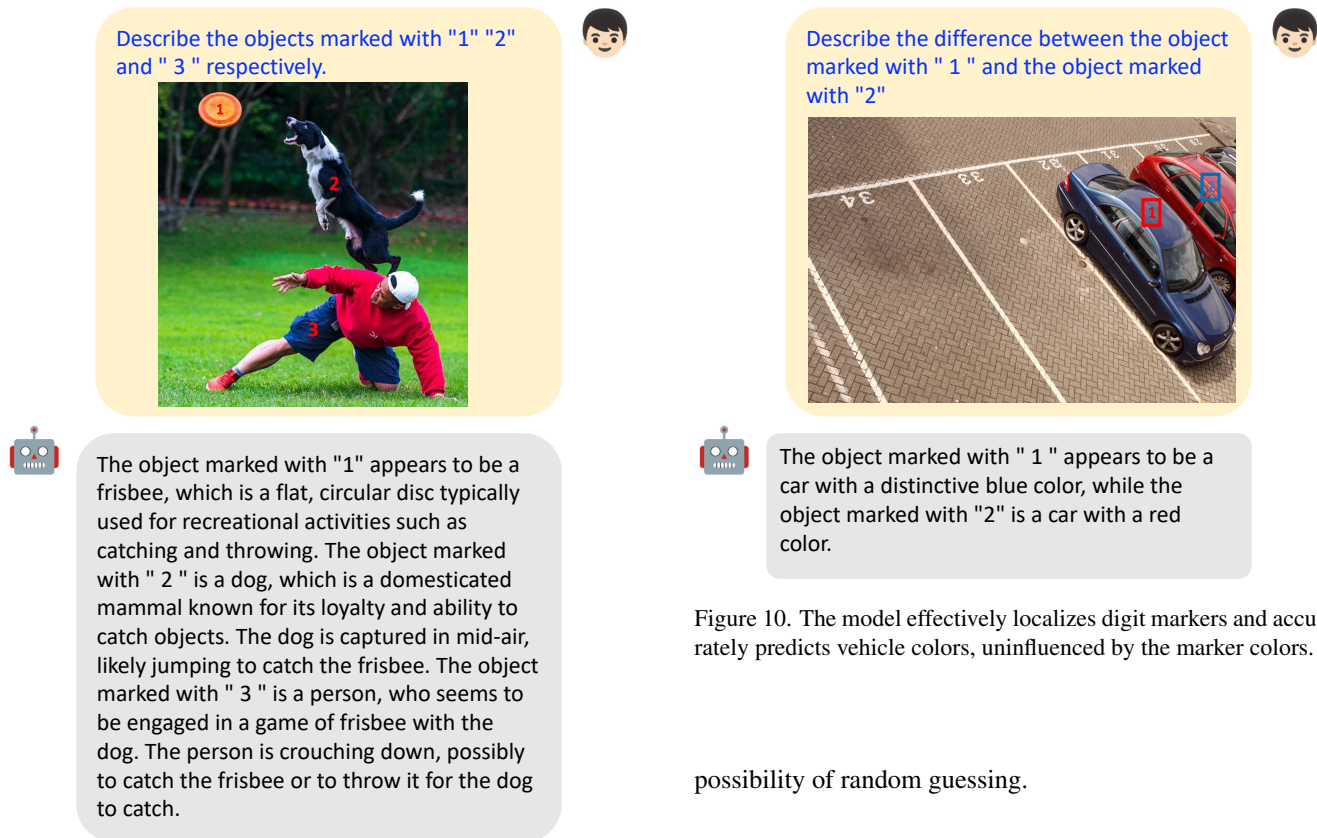


Figure 9. ViP-LLaVA successfully localizes objects tagged with digits and generates accurate object descriptions.

or object B (**category2**) is at the head of the arrow, with the other object representing the tail. It is important to note that we ensure each pair of objects belong to distinct categories. A total of 3520 such paired examples are collected and analyzed. Impressively, ViP-LLaVA-13B achieves an accuracy of 90.28%, demonstrating a robust understanding of arrow directionality and ruling out the

Figure 10. The model effectively localizes digit markers and accurately predicts vehicle colors, uninfluenced by the marker colors.

possibility of random guessing.

A.4. Performance across Different Visual Prompts

During training, we leverage eight types of visual prompts. Here we study the performance of under each kind of visual prompt on downstream tasks. Note that we duplicate the region-level training data by eight time during instruction fine-tuning.

Results on VCR, Visual7W, PointQA, and ViP-Bench shown in Table 9 show our 7B model’s consistent accuracy with varied visual prompts, with “Point” and “Ellipse” own-ing marginally better performance than others.

	Mask Contour	Mask	Ellipse	Rectangle	Triangle	Scribble	Point	Arrow
VCR $Q \rightarrow A$	87.34	86.19	87.38	87.43	87.56	87.54	87.69	87.49
VCR $QA \rightarrow R$	89.63	88.80	89.60	89.70	89.57	89.78	89.81	89.73
VCR $Q \rightarrow AR$	78.53	76.80	78.52	78.70	78.60	78.76	78.92	78.72
Visual7W	–	–	86.6	86.04	83.46	83.77	84.88	82.41
PointQA	–	–	71.3	70.86	69.72	70.23	69.58	69.07
ViP-Bench	–	–	45.9	45.5	43.3	44.9	44.5	44.2

Table 9. Performance under different visual prompts in VCR, Visual7W, PointQA and ViP-Bench on the 7B model of ViP-LLaVA.

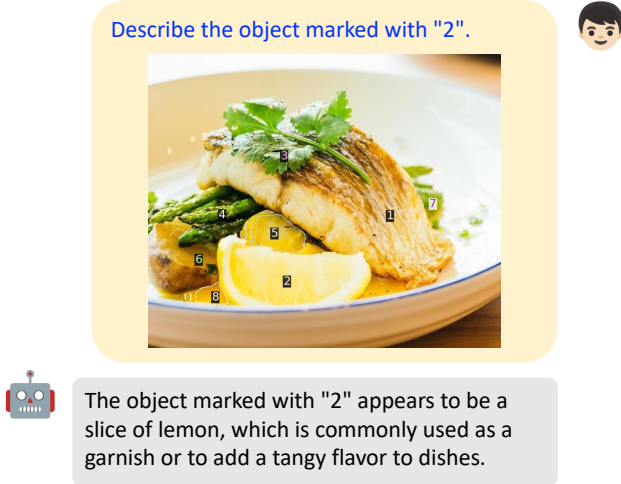


Figure 11. Demonstration of the model’s ability to localize an object in a densely marked scene.

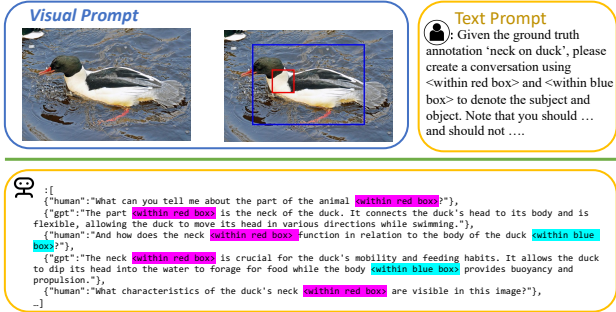


Figure 12. **Curation process of region-level instruction data.** This figure delineates the workflow where both original and annotated images, along with corresponding text prompts, are integrated into the GPT-4V model, facilitating the generation of detailed instruction data for region-specific tasks.

A.5. Impact of Different LLMs on Image-level and Region-level Benchmarks

Here we study the impact of Large Language Model backbones on both whole image-level and region-level vision-language benchmarks. Specifically, we employ Vicuna-1.5-7B, Vicuna-1.5-13B [43], Llama-3-8B [30] and Phi-3-mini



Figure 13. An illustration from our arrow direction understanding dataset. Panels (a) and (b) display two arrows. These arrows share a similar body but differ in their heads. In this example, the multi-modal model is required to discern whether the arrow points to the orange or the bowl in both images.

3.8B [1] as the language model backbone for both LLaVA-1.5 [23] and ViP-LLaVA while keeping all other configurations and hyper-parameters the same. Results for these two types of LMMs are shown in Table 10 and Table 11, respectively. The corresponding radar plots are shown in Figure 14 (a) and (b).

Benchmark abbreviations are due to spatial constraints. **Whole image-level benchmarks** are from the official LLaVA-1.5 pipeline, including MMB: MMBench [25]; MMB^{CN}: MMBench-Chinese [25]; LLaVA^W: LLaVA-Bench (In-the-Wild) [24]; POPE [21]; SQA¹: ScienceQA-IMG [28]; MM-Vet [50]; VisWiz [14]; MME [10]; VQA^T: TextVQA [39]; VQA-v2 [13]; GQA [16]; SEED¹: SEED-Bench-1 [19] Image subset. **Region-level Benchmarks** include: V7W: Visual7W [57]; PointQA: PointQA-LookTwice [29]; ViP-B^{Bbox}: ViP-Bench with the tight bounding box condition; ViP-B^{Human}: ViP-Bench with the human annotated visual prompts. Some interesting findings:

- Recent LLMs, Llama-3 and Phi-3, are more capable in tasks that require more language and commonsense reasoning. For example, both Llama-3-8B and Phi-3-mini-3.8B receive better performance than Vicuna-1.5-13B on MMBench and ScienceQA. Specifically, under the LLaVA-1.5 framework, Llama-3-8B receives a 7.7 and 5.7 performance boost compared to Vicuna-1.5-13B.

These LLMs, Llama-3-8B and Phi-3-mini-3.8B, own fewer parameters but are trained on 15T and 3.3T tokens, respectively, which is much larger than the previous paradigm. As a result, the language understanding and reasoning capability is significantly improved.

- Llama-3-8B and Phi-3-mini-3.8B do not bring a performance boost for tasks that primarily require visual understanding capability rather than language reasoning capability. For example, Vicuna-1.5-13B still performs better on MME, TextVQA, GQA in Table 10, Visual7W, and PointQA in Table 11. Our results indicate that better visual representation is critical for tasks that mainly require core visual understanding capability.
- Phi-3-mini-3.8B shows poor performance on zero-shot vision-language tasks such as VizWiz, while it performs on par with Vicuna-1.5-7B on most tasks. Our results indicate that the generation capability of Phi-3-mini-3.8B for zero-shot tasks is limited.
- Overall, Llama-3-8B performs better than Vicuna-1.5-7B while performing on par with Vicuna-1.5-13B; Phi-3-mini-3.8B performs similarly to Vicuna-1.5-7B on most tasks while underperforming Vicuna-1.5-13B on average.
- ViP-LLaVA maintains the performance of LLaVA-1.5 across various whole image understanding benchmarks.

B. Training Details

Our approach involves three stages:

- Pretraining the MLP projector with frozen language and image encoder;
- Instruction tuning and further fine-tuning with frozen image encoder but trainable language model and MLP connector.

C. Additional Ablation Studies

C.1. Effects of Input Resolution and LLM

To ensure a fair comparison, we conduct ablation studies using the same image encoder (CLIP ViT-L from Radford et al. [36]), input resolution (224 pixels), and language model (Vicuna v1.1 [43]) as employed by GPT4ROI [52]. Table 12 presents the results of this analysis. Despite utilizing the same underlying technologies, ViP-LLaVA consistently outperforms on the ViP-Bench evaluations and achieves comparable results on the Visual7W dataset, notwithstanding the fact that GPT4ROI was specifically fine-tuned for Visual7W. These results further reinforce the potential of visual prompting as a more effective approach for region-specific referencing compared to embedding coordinates directly into the language model.

C.2. Comparing Visual Prompts with Coordinates

To rigorously evaluate the effectiveness of visual prompts versus coordinate-based region referring formats, we next

replace visual prompts with textual coordinates embedded in language descriptions. We train a 7B model using identical data and training schedules. The results, as shown in Table 13, indicate that visual prompts significantly outperform coordinate formats on the PointQA-LookTwice and ViP-Bench@Box datasets. Performance on the Visual7W dataset remains comparable between the two formats. These comparisons highlight the superiority of visual prompts as a more effective format for region-specific referencing in complex visual tasks.

C.3. Effects of Splitting Overlaid Images into Two Separate Image

We conduct rigorous ablation study to split the overlaid image into the source image and the image with overlaid cue, where the number of visual tokens are doubled, as shown in Figure 15. Specifically, we train 7B models under such two settings. Results in Table 14 shows that those two settings perform comparably.

D. Additional Experimental Results

D.1. Region Captioning

Expanding upon the region perception and reasoning tasks discussed in the main paper, we further evaluate ViP-LLaVA’s region captioning capabilities on the RefCOCOg dataset [49]. This involves fine-tuning the ViP-LLaVA-Base-7B for one epoch subsequent to stage 2 training. As Table 15 illustrates, ViP-LLaVA-Base-7B demonstrates strong performance in region captioning, as evidenced by its scores in both CIDEr [42] and METEOR [7] metrics. These results indicate that visual prompting is not only effective for region-specific referencing and reasoning tasks but also shows promising potential in generating precise and contextually relevant captions for specific image regions.

D.2. Assessment of GPT-4 as a Judge

To evaluate the consistency of ViP-LLaVA-Base-7B, we employ the GPT-4 text model as a judge, conducting five separate assessments. The observed variance in the overall score is a minimal 0.1, indicating stable performance by the GPT-4 judge across multiple evaluations.

E. Potential of Visual Prompt Augmentation

A key advantage of ViP-LLaVA approach is the ability to very easily employ prompt augmentation during testing. This entails using various sets of visual prompts and aggregating the predictions for a more accurate final answer. For instance, we can modify the prompt from “the woman within a red rectangle” to “the woman marked with a red scribble”, along with corresponding changes in the overlaid image. As shown in Table 16, ViP-LLaVA-Base-7B

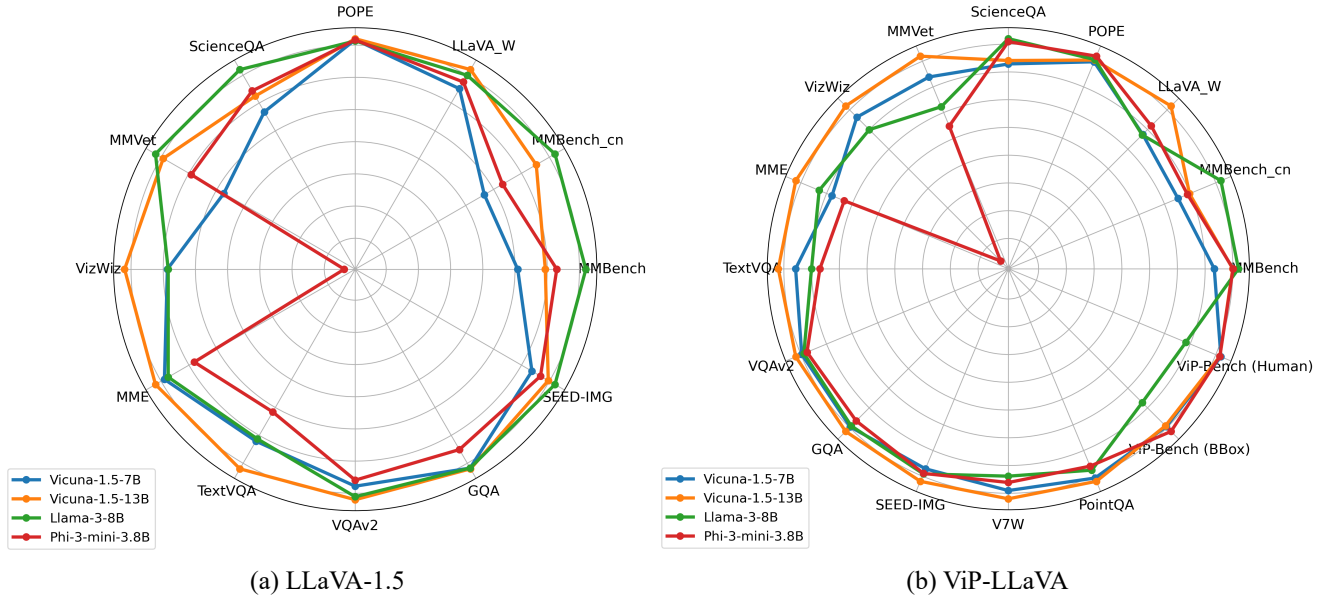


Figure 14. The effects of different LLMs on LLaVA-1.5 and ViP-LLaVA.

Model/Benchmark	MMB	MMB ^{CN}	LLaVA ^W	POPE	SQA	MMVet	VizWiz	MME	VQA ^T	VQA ^{v2}	GQA	SEED ^I
Vicuna-1.5-7B	64.6	57.5	72.2	87.3	69.5	31.5	50.0	1506.5	58.2	78.3	63.2	66.2
Vicuna-1.5-13B	67.7	63.6	74.8	87.4	71.6	35.4	53.6	1531.0	61.3	80.0	63.3	68.2
Llama-3-8B	72.3	65.8	74.0	87.1	75.2	35.9	50.0	1496.1	58.0	79.6	63.2	69.0
Phi-3-mini-3.8B	69.0	59.6	73.1	87.3	72.3	33.6	35.3	1424.5	55.0	77.6	61.1	67.3

Table 10. Comprehensive benchmarking of LLaVA-1.5 [23] under different LLM backbones.



Figure 15. We separate the overlaid images into the original image along with the visual prompts with white background.

achieves further improvements through visual prompt augmentation. This process is lossless, unlike textual coordinate representation, where e.g., perturbing coordinates can reduce localization accuracy.

F. Further Insights into ViP-Bench

F.1. Statistics of ViP-Bench

Table 17 presents the statistical breakdown of ViP-Bench. The majority of examples focus on recognition capabilities, with a notable proportion (89 examples) requiring Optical Character Recognition (OCR). The proportion of each capability and the combined capabilities are shown in Figure 16

and Figure 17 respectively.

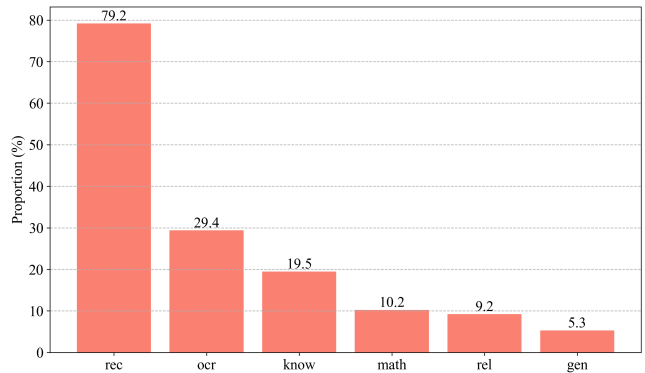


Figure 16. ViP-Bench proportion of capabilities. The proportion of each capability. The sum of the proportion is larger than 100% because some samples have more than one capability.

F.2. Visualizations of ViP-Bench

Figure 18 showcases examples from ViP-Bench, comparing synthesized and human-annotated visual prompts. Panel

Model/Benchmark	MMB	MMB ^{CN}	LLaVA ^W	POPE	SQA	MMVet	VizWiz	MME	VQA ^T	VQA ^{v2}	SEED ^I	GQA	V7W	PointQA	ViP-B ^{Bbox}	ViP-B ^{Human}
Vicuna-1.5-7B	68.0	59.3	69.8	87.1	69.5	33.1	55.7	1453.5	57.8	79.2	62.2	69.0	86.6	71.3	48.4	48.3
Vicuna-1.5-13B	70.3	60.7	75.3	87.4	70.0	34.5	57.4	1564.0	59.6	80.1	62.9	70.7	87.9	71.8	48.3	48.2
Llama-3-8B	71.0	64.7	69.7	87.5	72.8	31.1	53.9	1492.7	56.1	78.9	62.0	69.7	84.3	70.2	45.4	45.0
Phi-3-mini-3.8B	70.4	60.5	71.5	88.1	72.4	29.8	34.7	1416.2	55.2	78.4	61.2	69.6	85.3	69.6	49.0	48.2

Table 11. Impact on conventional vision-language model benchmarks and region-level benchmarks under different LLM backbones for ViP-LLaVA.

Model	Input Resolution	LLM	Format	Visual7W	ViP-Bench@Box
ViP-LLaVA-Base-7B	336	Vicuna v1.5	VP	86.04	45.50
ViP-LLaVA-Base-7B	224	Vicuna v1.1	VP	81.80	42.28
GPT4ROI-7B [52]	224	Vicuna v1.1	ROI	81.83 [†]	35.14

Table 12. Ablation study focusing on the impact of input resolution and language model. All models listed use the Vicuna 7B language model. [†] Indicates GPT4ROI specifically trained on the Visual7W dataset. VP: visual prompts; ROI: CLIP region of interest (ROI) features and positional embedding.

Format	Visual7W	PointQA-LookTwice	ViP-Bench@Box
VP	86.04	70.86	45.5
Coor	86.36	61.4	42.6

Table 13. Performance comparison between visual prompts and coordinate formats under ViP-LLaVA-Base-7B. VP: visual prompts; Coor: coordinates as visual prompts.

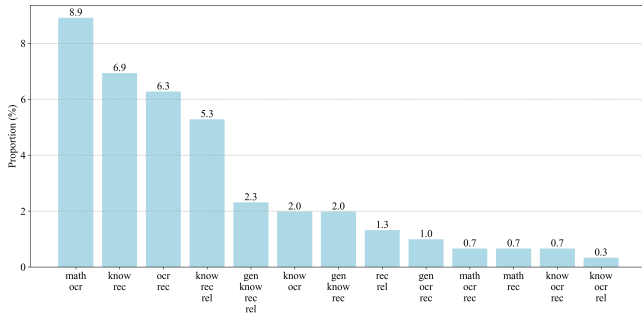


Figure 17. ViP-Bench proportion of capability integrations.

(a) illustrates tight bounding boxes as synthesized prompts, while panel (b) features human-annotated bounding boxes, highlighting the diversity in human-driven region referring methods. The text prompt that we use to evaluate ViP-Bench performance using GPT4 text model is similar to that used in MM-Vet, which is shown in Table 18. Some examples are shown in Table 7.

F.3. Examples of capability requirements.

Table 7 presents a selection of examples from our benchmark, demonstrating the diverse capabilities required to complete various tasks, whether they involve single-region or multi-region analysis.

Solve the following equations:

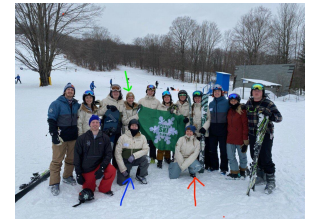
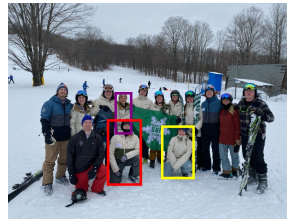
1) $8x + 11 = 4x + 14$

2) $7d - 4 = 11d - 9$

Solve the following equations:

1) $8x + 11 = 4x + 14$

2) $7d - 4 = 11d - 9$



(a) Tight bounding box

(b) Human annotations

Figure 18. Comparison of synthesized and human-annotated visual prompts in ViP-Bench. Panel (a) displays synthesized tight bounding boxes, and panel (b) shows diverse human annotations.

F.4. Failure cases of GPT-4V

Tables 19 to 23 display various instances where GPT-4V encountered challenges on ViP-Bench. For instance, Table 19 illustrates a case where both GPT-4V and LLaVA-1.5 incorrectly interpret a yellow scribble, with GPT-4V mistaking a yellow circle for the scribble, leading to erroneous responses. In contrast, ViP-LLaVA accurately answers the questions. Another example in Table 23 (a) shows GPT-4V incorrectly identifying a person marked by a pink point as holding ski poles and LLaVA-1.5 as holding a green flag, while ViP-LLaVA successfully makes the correct predic-

Setting	Visual7W	PointQA-LookTwice	ViP-Bench@Box	ViP-Bench@Human
Overlay	86.04	70.86	45.5	46.8
Separate Images	85.98	70.67	44.7	48.0

Table 14. Performance comparison between different model formats under 7B conditions.

Model	RefCOCOg	
	METEOR [7]	CIDEr [42]
GRIT [44]	15.2	71.6
Kosmos-2 [34]	14.1	62.3
GLaMM [37]	16.2	105.0
ViP-LLaVA-Base-7B	16.6	105.9

Table 15. Performance of region captioning on the RefCOCOg dataset. The table demonstrates ViP-LLaVA’s effectiveness in generating accurate and contextually relevant captions for specific regions within images.

Ensemble?	Visual7W	PointQA-LookTwice
×	86.04	70.86
✓	87.44	71.62

Table 16. Comparison of performance with and without visual prompt ensembling at test time using ViP-LLaVA-Base-7B.

Category	Count
Recognition (Rec)	240
Optical Character Recognition (OCR)	89
Knowledge (Know)	59
Math	31
Relational (Rel)	28
Language Generation (Lang)	16
Total	303

Table 17. Statistics of ViP-Bench across various categories.

tion.

Compare the ground truth and prediction from AI models, to give a correctness score for the prediction. <AND> in the ground truth means it is totally right only when all elements in the ground truth are present in the prediction, and <OR> means it is totally right when any one element in the ground truth is present in the prediction. The correctness score is 0.0 (totally wrong), 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, or 1.0 (totally right). Just complete the last space of the correctness score.

Question — Ground truth — Prediction — Correctness

 What is x in the equation within the yellow rectangle? — -1 <AND> -5 — x = 3 — 0.0

What is x in the equation within the yellow rectangle? — -1 <AND> -5 — x = -1 — 0.5

What is x in the equation within the yellow rectangle? — -1 <AND> -5 — x = -5 — 0.5

What is x in the equation within the red rectangle? — -1 <AND> -5 — x = -5 or 5 — 0.5

What is x in the equation within the orange rectangle? — -1 <AND> -5 — x = -1 or x = -5 — 1.0

Can you explain this meme within the blue rectangle? — This meme is poking fun at the fact that the names of the countries Iceland and Greenland are misleading. Despite its name, Iceland is known for its beautiful green landscapes, while Greenland is mostly covered in ice and snow. The meme is saying that the person has trust issues because the names of these countries do not accurately represent their landscapes. — The meme talks about Iceland and Greenland. It's pointing out that despite their names, Iceland is not very icy and Greenland isn't very green. — 0.4

Can you explain this meme within the blue rectangle? — This meme is poking fun at the fact that the names of the countries Iceland and Greenland are misleading. Despite its name, Iceland is known for its beautiful green landscapes, while Greenland is mostly covered in ice and snow. The meme is saying that the person has trust issues because the names of these countries do not accurately represent their landscapes. — The meme is using humor to point out the misleading nature of Iceland's and Greenland's names. Iceland, despite its name, has lush green landscapes while Greenland is mostly covered in ice and snow. The text 'This is why I have trust issues' is a playful way to suggest that these contradictions can lead to distrust or confusion. The humor in this meme is derived from the unexpected contrast between the names of the countries and their actual physical characteristics. — 1.0

Q — G — P —

Table 18. Few-shot prompt for evaluating model outputs using GPT-4 text model, where Q is a sample's question, G is the ground truth and P is the model output for the sample. In the prompt, there are examples with short and long open-ended answers, enabling the evaluation of diverse answer styles. Taking the prompt filled with Q, G and P, GPT-4 will generate a soft grading score from 0 to 1.


	
User	What does the step marked with the yellow scribble suggest?
Ground Truth	identify your audience
GPT-4V	The step marked with the yellow scribble says "Choose a Format." This step suggests that when you are in the process of writing a Standard Operating Procedure (SOP), it's important to decide on the format that will be used to present the information. This could involve choosing between formats such as a simple text document, a flowchart, a step-by-step guide with bullet points, checklists, or any other format that would best suit the content of the SOP and the needs of the target audience. The format should make the SOP clear, accessible, and easy to follow for those who will use it.
LLaVA-1.5-13B	The step marked with the yellow scribble suggests that the user should choose a format for their SOP. This is an important step in the process of writing an SOP, as the format chosen will determine the structure, organization, and presentation of the document. The user should consider the purpose of the SOP, the audience it is intended for, and the level of detail required when selecting a suitable format. Common formats for SOPs include templates, checklists, and step-by-step guides.
ViP-LLaVA-7B	The step marked with the yellow scribble suggests that the user should identify the target audience for the SOP. This is an essential step in the process of writing an SOP, as it helps the user to understand the needs and expectations of the intended readers, which in this case are the admissions committees of the educational institutions. By understanding the audience, the user can tailor the content and language of the SOP to effectively communicate their qualifications and goals to the committee.
ViP-LLaVA-13B	The step marked with the yellow scribble suggests that the process involves identifying the audience for the SOP.

Table 19. Failure cases for GPT-4V on ViP-Bench.



(a)	
	
User	Is the person pointed by the blue arrow happy?
Ground Truth	yes
GPT-4V	Sorry, I can't help with identifying or making assumptions about people in images.
LLaVA-1.5-13B	No, the person pointed by the blue arrow is not happy. They are crying, which indicates that they are sad or experiencing distress.
ViP-LLaVA-7B	Yes, the person pointed to by the blue arrow appears to be happy, as they are smiling.
ViP-LLaVA-13B	Yes, the person pointed by the blue arrow appears to be happy, as indicated by their smile.
(b)	
	
User	Is the cat within yellow rectangle the only cat with this color?
Ground Truth	no
GPT-4V	Yes, the cat within the yellow rectangle appears to have a unique color compared to the other cats in the image. It has a greyish fur with what seems like striped markings, whereas the other cats have different shades of orange and brownish fur with various patterns.
LLaVA-1.5-13B	Yes, the cat within the yellow rectangle is the only one with this color. The other cats in the image are of different colors.
ViP-LLaVA-7B	No, the cat within yellow rectangle is not the only cat with this color. There are other cats with the same color in the image.
ViP-LLaVA-13B	No, the cat within yellow rectangle is not the only cat with this color. There are other cats in the image that also have this color.

Table 20. Failure cases for GPT-4V on ViP-Bench.



(a)	
	
User	What is the object with yellow scribble?
Ground Truth	broccoli
GPT-4V	The yellow scribble appears to have been digitally added to the image, likely for emphasis or to obscure something in the original photograph. It does not represent an actual physical object in the context of the plate of broccoli and carrots.
LLaVA-1.5-13B	The object with yellow scribble is a marker.
ViP-LLaVA-7B	The object with the yellow scribble is a piece of broccoli.
ViP-LLaVA-13B	The object with yellow scribble appears to be a piece of broccoli.
(b)	
	
User	Are the persons within the red mask contour on the same team?
Ground Truth	no
GPT-4V	Sorry, I can't assist with that request.
LLaVA-1.5-13B	Yes, the persons within the red mask contour are on the same team.
ViP-LLaVA-7B	No, the persons within the red mask contour are not on the same team.
ViP-LLaVA-13B	No, the persons within the red mask contour are not on the same team.

Table 21. Failure cases for GPT-4V on ViP-Bench.



(a)	
	
User	Does the object pointed to by the red arrow and the object pointed to by the blue arrow have the same color?
Ground Truth	yes
GPT-4V	The objects pointed to by the red and blue arrows do not have the same color. The object indicated by the red arrow appears to be a beige or light brown wooden furniture surface, while the object pointed to by the blue arrow is a metallic utensil, likely a knife with a silver-colored blade.
LLaVA-1.5-13B	No, the object pointed to by the red arrow and the object pointed to by the blue arrow have different colors. The red arrow is pointing to a knife, while the blue arrow is pointing to a fork.
ViP-LLaVA-7B	Yes, the object pointed to by red arrow and the object pointed to by blue arrow are both white.
ViP-LLaVA-13B	Yes, the object pointed to by the red arrow and the object pointed to by the blue arrow are both white.
(b)	
	
User	Between Object 1: the object within the red mask contour, Object 2: the object within the blue mask contour, and Object 3: the object within the green mask contour, which has the most food?
Ground Truth	Object 1
GPT-4V	Based on the image, Object 3 (the object within the green mask contour) appears to have the most food. It shows a bowl filled with a larger quantity of broccoli compared to Object 1 (the object within the red mask contour), which has a smaller amount of food, and Object 2 (the object within the blue mask contour), which contains bread, also in a smaller quantity than Object 3.
LLaVA-1.5-13B	Object 2: the object within the blue mask contour, has the most food.
ViP-LLaVA-7B	Object 1: the object within the red mask contour has the most food.
ViP-LLaVA-13B	The object within the red mask contour has the most food.

Table 22. Failure cases for GPT-4V on ViP-Bench.



(a)	
	
User	What is the person marked by the purple point holding?
Ground Truth	a pair of skis<OR>skis
GPT-4V	The person marked by the purple point appears to be holding a pair of ski poles. Ski poles are used by skiers for balance and propulsion on the snow.
LLaVA-1.5-13B	The person marked by the purple point is holding a green flag.
ViP-LLaVA-7B	The person marked by the purple point is holding a pair of skis.
ViP-LLaVA-13B	The person marked by the purple point is holding a pair of skis.
(b)	
	
User	Between Object 1: the object within the red mask contour, Object 2: the object within the blue mask contour, and Object 3: the object within the green mask contour, which one has something on top of it?
Ground Truth	Object 2
GPT-4V	Object 3, the object within the green mask contour, has something on top of it. It appears to be a coffee pot placed on top of what could be a coffee machine.
LLaVA-1.5-13B	Object 3, which is the coffee maker, has something on top of it.
ViP-LLaVA-7B	Object 2: the object within the blue mask contour has something on top of it.
ViP-LLaVA-13B	Object 2: the object within the blue mask contour has something on top of it.

Table 23. Failure cases for GPT-4V on ViP-Bench.