

Expand VSR Benchmark for VLLM to Expertize in Spatial Rules

Peijin Xie¹, Lin Sun², Bingquan Liu^{1*}, Dexin Wang², Xiangzheng Zhang², Chengjie Sun¹, Jiajia Zhang³

¹Faculty of Computing, Harbin Institute of Technology, Harbin, 150001, China,

²360 Search Department, Beijing, 100020, China,

³School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, 518055, China,

xpj@stu.hit.edu.cn, {liubq,sunchengjie,zhangjiajia}@hit.edu.cn,

{sunlin1, wangdexin, zhangxiangzheng}@360.cn

Abstract

Distinguishing spatial relations is a basic part of human cognition which requires fine-grained perception on cross-instance. Although benchmarks like MME, MMBench and SEED comprehensively have evaluated various capabilities which already include visual spatial reasoning(VSR). There is still a lack of sufficient quantity and quality evaluation and optimization datasets for Vision Large Language Models(VLLMs) specifically targeting visual positional reasoning. To handle this, we first diagnosed current VLLMs with the VSR dataset and proposed a unified test set. We found current VLLMs to exhibit a contradiction of over-sensitivity to language instructions and under-sensitivity to visual positional information. By expanding the original benchmark from two aspects of tuning data and model structure, we mitigated this phenomenon. To our knowledge, we expanded spatially positioned image data controllably using diffusion models for the first time and integrated original visual encoding(CLIP) with other 3 powerful visual encoders(SigLIP, SAM and DINO). After conducting combination experiments on scaling data and models, we obtained a VLLM VSR Expert(VSRE) that not only generalizes better to different instructions but also accurately distinguishes differences in visual positional information. VSRE achieved over a 27% increase in accuracy on the VSR test set. It becomes a performant VLLM on the position reasoning of both the VSR dataset and relevant subsets of other evaluation benchmarks. We hope it will accelerate advancements in VLLM on VSR learning.

Extended version — <https://github.com/peijin360/vsre>

Introduction

Reasoning on spatial relations is a basic part of human cognition which requires fine-grained perception on cross-instance. Traditional classification benchmark VSR (Liu, Emerson, and Collier 2023) has proposed a controlled probing dataset testing vision language models’ capabilities of discrimination on spatial relations with natural image-text pairs. But, as Large Language Models (LLMs) become research hotspots, traditional tasks have been incorporated

into conversational QA scenarios. The new VSR task requires the model to not only accurately recognize visual positional information but also follow instructions to correctly answer questions. Although models have evolved from traditional models with classification heads (ViLT, VisualBERT, and LXMERT) to VLLMs (LLaVA, BLIP2, Qwen-VL, and GPT-4), a unified evaluation method and effective optimization approach for VLLM toward VSR is still lacking.

VLLM is the advanced version of LLM, which is enhanced to process and interpret multi-modal data. They leverage a powerful LLM as their cognitive engine to handle various tasks. Equipped with visual tokens, LLMs can perceive rich visual information and perform complex reasoning like VSR. However, most VLLMs often fall into the trap of hallucination problems. During the model’s hallucination evaluation, the spatial relationships are precisely the type of relational hallucinations that are more challenging to object and attribute hallucinations (Liu et al. 2024c). Therefore, filling the blank of the evaluation and optimization of VLLM on VSR is of great significance for addressing the issue of hallucinations on relation.

We first re-evaluated the VLLM from scratch and diagnosed issues of inconsistent performance, hypersensitivity on text prompts, insensitivity on vision information and answer bias. We found the absence of a unified instruction test has resulted in significant variance in model performance and the instability of instruction-following hinders the evaluation and optimization of VLLM’s VSR capabilities. Therefore, we proposed a unified instruction test set by expanding the VSR test set through both manual and GPT4-generated templates.

Similarly, we expand the training set with the same template pool. Trained with more diverse QA formats, the model achieves fundamentally better generalization in responding to VSR questions regardless of the question style. Considering a VSR expert requires more of the ability to distinguish and recognize visual spatial information rather than just correctly following instructions, we controllably augmented and repainted the visual training image to specific spatial relation concepts (like “in”, “on” and “under”) by diffusion model. The increased amount and diversity of image data strengthen the model’s comprehension of visual spatial details. Furthermore, we expand the vision encoder CLIP (Radford et al. 2021) with other hot encoders

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(SAM (Kirillov et al. 2023), DINOv2 (Oquab et al. 2024) and SigLIP (Zhai et al. 2023)) to obtain more vision perception ability that promotes model sensitivity to spatial information. After the experiments on scaling laws by data augmentation and ablation study on encoder combination, we obtained a VLLM VSR Expert(VSRE) that not only generalizes better to different text instructions but also accurately distinguishes differences in visual positional information. Our expansion achieved over 27% increase in accuracy on the VSR bench. The VSRE achieved the best performance on the position reasoning of both the VSR dataset and subsets of other evaluation benchmarks. Additionally, the phenomenon of answer bias has also been alleviated as the model’s attention shifts towards the visual aspects.

The contribution of this work can be summarized as follows:

1. We proposed the first unified VSR instruction test set and identified critical issues of inconsistent performance, hypersensitiveness on text, insensitivity on vision and answer bias.
2. We expanded the training data nearly 100 times by rewriting the text and repainting the image.
3. We expanded the single visual backbone to a merged visual encoder which enhanced the VLLM’s vision perception ability.
4. Our proposed VSRE overcomes the issues. It not only generalizes better to different text instructions but also accurately distinguishes differences in visual positional information.

Related Works

VSR Tasks and Datasets

The VSR dataset (Liu, Emerson, and Collier 2023) is a controlled dataset that explicitly evaluates traditional pretrained vision language models like VisualBert, LXMERT and Vilt. It contains more than 10k natural text-image pairs with 66 types of spatial relations which is still insufficient for instruction tuning for VLLM. During the image data selection from MSCOCO (Lin et al. 2014), they used a contrastive caption generation approach and manual second-round validation to guarantee clear positional relationship features and data balance. As each image-text pair sample contains a triplet of [*subject*(with bbox), *relation*, *object*(with bbox)], the VSR dataset has excellent extensibility under the QA scenario where VLLM fits.

Existing hot benchmarks like MME (Fu et al. 2024), MM-bench (Liu et al. 2023b), SEED (Li et al. 2023) test VLLMs’ various capabilities and also have integrated assessments of VSR capability. However, the amount of their relevant subsets, only makes them suitable for evaluating model capabilities but insufficient for tuning the VLLM.

Previous works such as InstructBlip (Dai et al. 2023), LLaVA (Liu et al. 2024a,b, 2023a) and Mm1 (McKinzie et al. 2024) transforms existing VQA benchmarks (Krishna et al. 2017; Goyal et al. 2017) into instruction tuning data, showing marked VLLM performance improvements. With

this inspiration, we take well-controlled and balanced original VSR triplets as seeds, expand them dozens of times to vision instruction data to fill in the gaps between traditional visual language models and VLLM.

Several VLLMs have posted their performance on VSR. Due to the lack of a unified test template, their result seemed inconsistent through reproduction. We specifically discuss this issue in the Section **Re-evaluate from Scratch**.

Merged Vision Encoder

Current works explored the benefits of multi-visual joint encoding across various perception and cognition tasks. Their ablation experiments identified the best combinations of popular vision features for specific tasks.

Mixture of Features (MoF) (Tong et al. 2024b) demonstrated that integrating vision self-supervised learning DINOv2 (Oquab et al. 2024) features with VLLMs can significantly enhance their visual grounding capabilities. MG-LLaVA (Zhao et al. 2024) enhances the model’s visual processing capabilities by incorporating a multi-granularity vision flow, which includes low-resolution, high-resolution, and object-centric features. Prism (Karamcheti et al. 2024) perform a head-to-head comparison between CLIP (Radford et al. 2021), SigLIP (Zhai et al. 2023), DINOv2, and a standard Vision Transformer pretrained for classification (on ImageNet-21K, finetuned on ImageNet-1K) and find that the backbones trained with vision-language contrastive objectives (i.e. CLIP and SigLIP) are significantly more performant than alternatives. Cambrian-1 (Tong et al. 2024a) also explored 20 vision encoders and their combinations. They conclude that High-res encoders greatly enhance performance on chart&vision-centric benchmarks. Combining multiple vision encoders, including Self Supervise Learning models, enhances VLLM performance across various benchmarks, particularly in vision-centric tasks.

Motivated by their findings, we summarized their high-performing backbones(including SigLIP, DINOv2 and SAM (Kirillov et al. 2023)) and followed their combination choice to expand our vision feature. After all, when it comes to intuitively understanding VSR tasks, precise visual detail perception and discrimination capabilities seemed more essential than complex semantic reasoning abilities. This necessitates that the model’s visual feature processing component be more sensitive to visual positional information and possess more specialized spatial encoding capabilities.

Controlled Image Generation

Diffusion models for controlled image generation offer a powerful and flexible approach to creating images that meet specific criteria. By incorporating conditioning information into the denoising process, models can generate highly controlled and precise outputs. They represent a powerful approach for image repainting, leveraging the systematic denoising process to generate high-quality, contextually consistent images. In this work, we utilized the most 3 popular applications to expand the image data as follows:

Text to Image generates an image from a text description. The denoising process is guided by the text, and once the

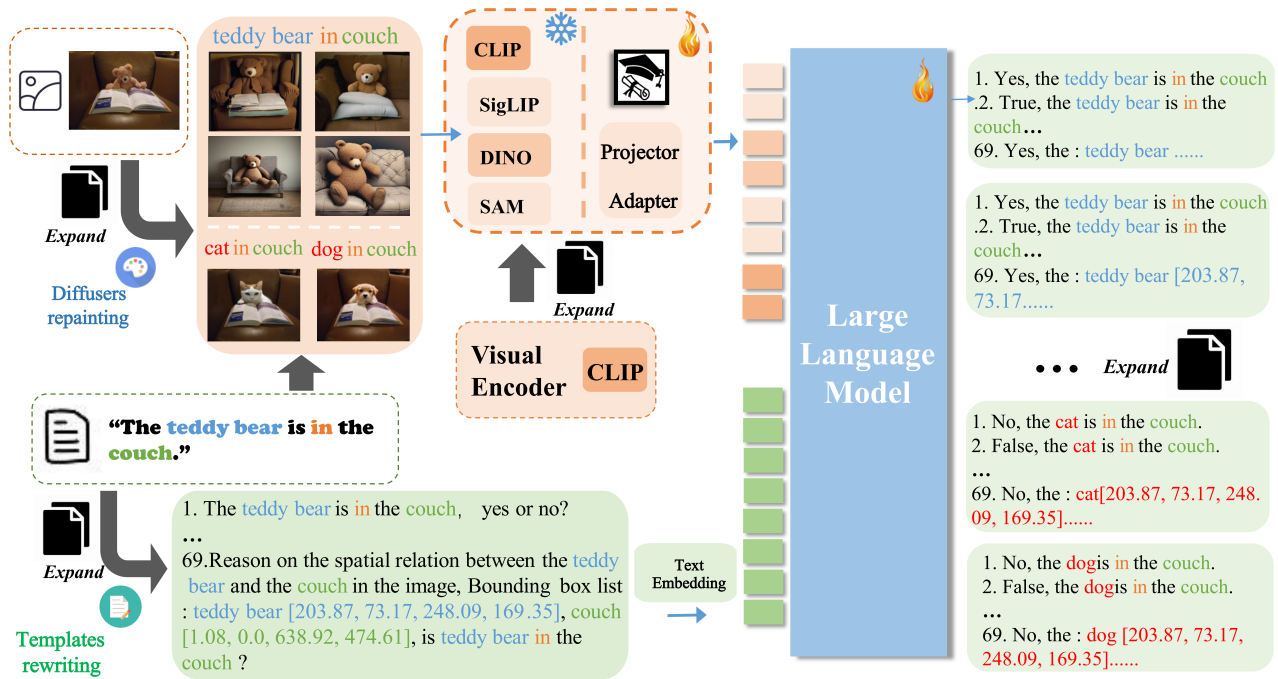


Figure 1: The overall expansion method through the training process. On the text branch, questions and answers are rewritten by templates in green blocks. On the image branch, the image inputs are repainted by the diffusion model through *image to image*, *text to image* and *image inpainting* 3 methods in respectively 3 rows in the left orange block. And the expansion on vision encoder to a powerful merged one is shown in the middle orange block with dashed box.

denoising process ends after a predetermined time steps, the image representation is decoded into an image.

Image-Text to Image is similar to text-to-image, but in addition to a prompt, an initial image is encoded to latent space then the noise is added to it. Then the model takes a prompt and the noisy latent image, predicts the noise, and removes the predicted noise from the initial latent image.

Image Inpainting replaces or edits specific areas of an image. This makes it a useful tool to replace an image area with something entirely new. It relies on a mask to determine which regions of an image to fill in; the area to inpaint is represented by white pixels and the area to keep is represented by black. The white pixels are filled in by the prompt.

Re-evaluation from Scratch

In this section, we re-evaluate the performance of popular VLLMs on the VSR test from scratch and identify the problem of inconsistent performance, hypersensitiveness on text, insensitivity on vision and answer bias that hinder the evaluation. Due to page limitations, we provide a brief overview of our re-evaluation process as follows, more details are shown in Appendix Material Section 1.

We first provide the most comprehensive summary including over 120 attempts across more than 30 models. During the summary constitution, We identified inconsistencies in model performance, such as the model LLaVA1.5 having nearly a 20% lower accuracy in MiniGPTv2’s reproduction(51%) (Chen et al. 2023) compared to the result in Prism(71%) (Karamcheti et al. 2024).

Then we selected the most commonly used VLLM architecture, LLaVA1.5, for prompt engineering experiments with 69 templates. We found that hypersensitiveness on text prompts caused the inconsistency in performance. When asked about positional information, the model’s accuracy is significantly affected by factors such as the questioning format, the insertion or deletion of specific phrases, and the order of certain words. Similarly, through case studies, we found that insensitive visual features struggle to distinguish positional categories.

Furthermore, the model exhibited severe response biases for most template answers. Questions with “yes” answers have a significantly higher accuracy rate compared to those with “no”. We suspect that may be due to the co-occurrence of subject and object concepts in both question text and images. This co-occurrence phenomenon may confuse the model, leading it to hastily provide “yes” answers based on co-occurrence rather than focusing on the actual visual spatial relationship.

To conclude, the over-sensitivity to language instructions caused large variance and inconsistency in model performance, while under-sensitivity to visual information led to insufficient perception of spatial relations, affecting answer accuracy and introducing response bias.

Expansion for Spatial Expert

To overcome the issues and obtain a Spatial Expert, we arrange our expansion method through VLLM’s pertaining and instruction tuning pipeline in Figure 1.

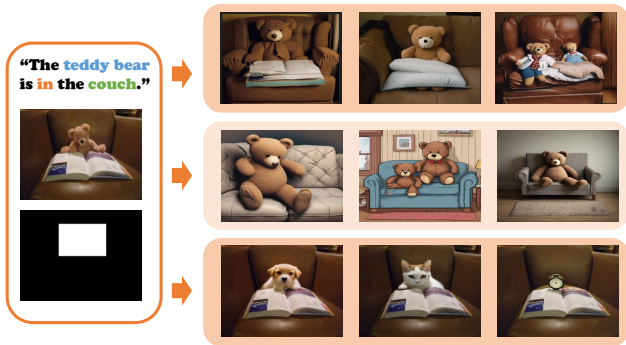


Figure 2: Examples of 3 settings of image-to-image(first row), text-to-image(middle row), and inpainting(last row) through the repainting process with the original image-text pair and mask inputs on the left.

Expansion on Text Data

Inspired by the success of visual instruction tuning, we prepare the test training data from the original VSR training [subject, relation, object] triplets extracted by spaCy¹. We hope that after instruction fine-tuning (IFT) with augmented text data, the model, having been exposed to more diverse forms of Q&A, can reduce its sensitivity to question text prompts and exhibit higher generalization on instruction following when faced with various questioning formats. We expand the text QA data through both manual and GPT-4o generated prompt templates. To be exact, we select the top 30 templates manually in the re-evaluation process and append 20 GPT4-o generated ones.

Expansion on Image Data

To enhance the perception of position on VLLM, we provide them with dozens of times more image inputs with spatial relation concepts than ever. Motivated by the success on LLaVA pretrained on GPT4 generated captions, inversely, we freeze the text captions as prompts and utilize image generation diffuser SDXL (Podell et al. 2023) to generate images. More details and samples of the generation are posted in the Appendix Material Section 2.

The general image generation under prompt control ensures the display of specific positional relationship information between two entities, such as "The teddy bear is in the couch". We believe that an efficient principle for data augmentation is to generate images with as much diversity as possible while ensuring the original positional relationship semantics are maintained. Therefore, we employed the most three general settings to sequentially increase diversity.

As shown in Figure 2, **Image-to-Image** repaint the picture slightly in the first row which also maintained the consistency of the color on the background and the white region in front of the bear. The differences might not be apparent after visual encoding.

Therefore, we armed our augmentation with **Text-to-Image** and **Image-Inpainting** strategic. Text-to-Image gen-

¹<https://spacy.io/>

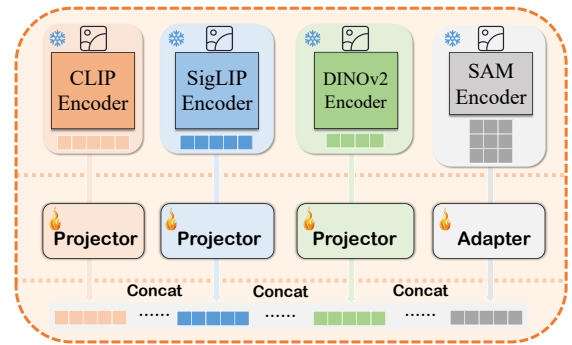


Figure 3: Illustration of the Merged Vision Encoder that concatenate multiple visual features aligned by projector or adapter respectively.

erated a more diverse set of images including variations in background and couch color and image style in the second row of Figure2. And in the last row, we utilized the bounding box as mask to inpaint the photo to make alter on *subject* and *object*. We replaced relatively small items with similar concepts like the bear with a dog, a cat and a clock. The revision helps prevent the model from learning biases on the distribution of *subject* and *object* combinations in the text.

Expansion on Vision Encoder

With sufficient VSR text and image data support, we incorporated the most used visual backbone CLIP² with SigLIP³, DINOv2⁴ and SAM⁵ to fully explore the potential of combining visual features. The language-guided contrastive model CLIP and SigLIP benefit from the massive scale of noisy web image-text data. But self-supervised encoders DINOv2 and segmentation SAM may detect more fine-grained visual details which may profit the spatial reasoning process. We then merge the pretrained backbones on the shelf to obtain a superior vision encoder focusing on visual spatial relations. Specifically, as shown in Figure3, we design projectors and adapter to align visual tokens, concatenate these tokens along the feature dimension following MOF, MG-LLaVA and Cambrian. Additional designs of the projector and adapter are shown in Appendix Section 3.

Experiment

In this section, we illustrate the details of the dataset creation and training&inference arrangement. As for the baseline model, we adopt the most extensively adopted VLLM architecture LLaVA1.5 to verify our expansion methods. Similar to its training process, we also arrange our optimized experiment into 2 stages: (1) Pertaining stage to enhance the perception ability of visual spatial relation and understanding of positional text concepts. (2) Instruction fine-tune (IFT) stage to enhance instruction following ability then

²<https://huggingface.co/openai/clip-vit-large-patch14-336>

³<https://huggingface.co/google/siglip-so400m-patch14-384>

⁴<https://huggingface.co/facebook/dinov2-base>

⁵<https://huggingface.co/facebook/sam-vit-base>

reduce the sensitivity to question text prompts and obtain higher generalization towards variant question formats.

Datasets

Testing Datasets To maintain consistency with previous evaluations, we used the VSR zero-shot test sets as the basis. We organized our test data into two sets:

(1)*Test-G* random sampled a prompt from the 50 templates pool for each triplet to evaluate instruction-following generalization ability during spatial reasoning.

(2)*Test-S* froze the template to the specific one (*[caption]*, True or false.) which performs the best and is the simplest during the re-evaluation process. The split aims to test the model’s VSR capability on visual under its most proficient question-asking format.

Training Datasets Excluding the test set, we collected more than 10k triplets with images from the original VSR dataset as a seed. Then we expand it several dozens of times into pre-train and IFT data as follow:(more details in Appendix Section 4)

- Pre-training data: Under the three settings with a ratio of 5:3:2, we repainted the original images, expanding the quantity 20 to 100 times the original amount. We label the set as “pre-100k” for 100k pre-training data and “pre-500k” for 500k.
- IFT data: We used general 50 prompt templates (30 manual and 20 GPT4-generated) to expand the 11k triplet data nearly 50 times to 500k, then name it as “turn-g 500k”. Note that “turn-s 11k” and “turn-g 11k” are unexpanded IFT data for comparison. “turn-g 11k” used a randomly selected temple for each triplet and “turn-s 11k” used the specific template same as *Test-S*.

Training and Inference

To ensure fairness, we seed everything before the training and inference process. Throughout the entire training process, we froze all visual backbones. During the pre-training stage, we froze the LLM and only trained the adapter layers to encourage the model to learn better visual encoding of spatial details. In the fine-tuning stage, we unfroze the LLM, allowing it to participate in the training to enhance the model’s instruction-following capability. Finally, during the inference stage, we limited the length of the model’s responses to only one new word. This word was used as the answer to binary questions; for example, words like “True”, “Yes” or “A” were considered positive predictions, while “False”, “No” or “B” were considered negative. Any other responses were directly judged as incorrect.

Result and Analysis

Scaling on Data

We validated the effectiveness of our expanded data on LLaVA, further replaced the LLM and VLLM, the results showed that our data expansion adapted to various models.

Table 1 shows the result of LLaVA1.5 7B and 13B on scaling training data experiment. We pre-trained and tuned the existing projector(adapter) and LLM with their weights

Pretrain Adapter	IFT Adapter+LLM	acc 7B <i>Test-G/Test-S</i>	acc 13B <i>Test-G/Test-S</i>
-	-	54.3 / 65.3	57.7 / 68.4
-	turn-g 11k	-	-
turn-g 11k	-	57.3 / 65.7	59.2 / 68.2
-	turn-s 11k	-	-
turn-s 11k	-	55.1 / 67.9	56.7 / 70.1
-	turn-g 500k	58.3 / 69.5	62.5 / 71.4
pre-100k	turn-g 500k	61.7 / 71.0	63.2 / 73.7
pre-200k	turn-g 500k	64.1 / 73.3	65.8 / 74.9
pre-300k	turn-g 500k	65.6 / 74.1	69.7 / 75.5
pre-400k	turn-g 500k	66.7 / 73.5	70.3 / 75.7
pre-500k	turn-g 500k	66.2 / 73.6	70.2 / 75.6
pre-400k ₁ turn-s 11k ₃	turn-g 500k ₂	66.4 / 74.7	70.1 / 76.6

Table 1: Result of LLaVA1.5 7B and 13B on scaling training data experiment. We post the Test-G and Test-S accuracy (split through “/”) by pretrained the adapter with data of the first column and instruct fine turning(IFT) both the adapter and LLM with the second column data.

on-the-shelf using the augmented data. We found that during pre-training, adjusting the projector(adapter) with even just 11k of IFT data can lead to an increase in accuracy, but model failed to respond after tuning LLM by tiny 11k data. Therefore, we added the augmented 500k tune-g data and gradually increased the pre-train samples. First, after training with the turn-g 500k IFT samples, the 7B model improved by 4.0% on the *Test-G* set, and 13B model improved by 4.8%. Both versions achieved certain improvements on the *Test-S*. As the amount of pre-training data increased, the models continued to achieve better results. Although the accuracy gains became less noticeable when the data size increased beyond 400k, the 13B model ultimately achieved scores of 70.3 and 75.7 on each test. This indicates that data expansion not only helps to enhance the model’s generalization ability to answer various text questions but also improves the model’s ability to discern positional information. Notably, in the last row, we added tuning data turns-s 11k at the end, meaning we first pre-trained with pre-400k, then tuned with turn-g 500k, and finally pre-trained again with turn-s (the training sequence is marked with subscripts number in the table). As a result, we found that the model’s performance on the *Test-S* further improved to 76.6. This demonstrates that supplementing with relevant data has great potential for optimizing model performance on VSR.

Table 2 shows the result of scaling data across other hot-spot LLM and VLLM. We replaced the tuned LLM in LLaVA with other LLMs and randomly initialized the adapter for retraining. Also, we combined the pretrain and the tune set to fine-tune other VLLMs together. The results indicate that although the performance was not as good as continuing training on the trained LLaVA, our data augmentation method significantly improved other LLMs and VLLMs on VSR task. The accuracy of the model’s responses improved consistently in both randomly asked general questions and fixed-format questions.

LLM / VLLM	zero-shot	turn-g 500k	+pre-100k	+pre-200k	+pre-300k	+pre-400k	+pre-500k
vicuna 7B	-	56.4 / 64.0	58.3 / 67.2	60.2 / 68.2	61.6 / 70.1	63.1 / 72.6	63.4 / 72.9
vicuna 13B	-	59.8 / 67.3	62.4 / 69.1	64.7 / 70.8	65.7 / 73.5	68.7 / 74.2	69.2 / 74.2
LLAMA2 7B	-	57.1 / 61.7	57.8 / 64.5	61.3 / 67.3	62.1 / 68.2	63.3 / 69.9	62.4 / 70.1
LLAMA2 13B	-	60.9 / 63.2	61.8 / 67.8	65.5 / 70.6	67.0 / 72.4	69.2 / 73.4	68.8 / 74.4
LLAMA3 8B	-	61.5 / 65.4	62.6 / 68.4	65.8 / 70.5	68.5 / 72.9	70.0 / 74.1	70.0 / 74.3
Qwen-VL 7B	57.8 / 62.7	63.4 / 68.2	65.8 / 69.2	66.1 / 71.7	67.0 / 72.2	67.9 / 73.6	68.2 / 73.6
BILP2 (FlanT5XXL)	59.3 / 66.5	64.2 / 69.6	66.3 / 69.7	68.0 / 71.1	69.3 / 72.5	70.4 / 73.9	70.2 / 74.2
InstructBLIP (FlanT5XXL)	54.4 / 63.1	59.0 / 66.2	62.5 / 67.8	64.3 / 69.7	66.9 / 71.8	67.3 / 71.9	68.2 / 72.3

Table 2: Result of scaling data across other hot-spot LLM and VLLM on Test-G and Test-S (split through “/”). The column names represent the data used for training, in sequence of first 3 as: no data for “zero-shot”, only 500k turn-g data for “turn-g 500k”, turn-g 500k plus pre-100k for “+pre-100k”. For LLMs, we randomly initialized the adapter weights and sequentially used the corresponding data to perform the pre-train and fine-tune processes. And for VLLMs, we combined the tuning and pre-train data sequentially for IFT.

Scaling on Model

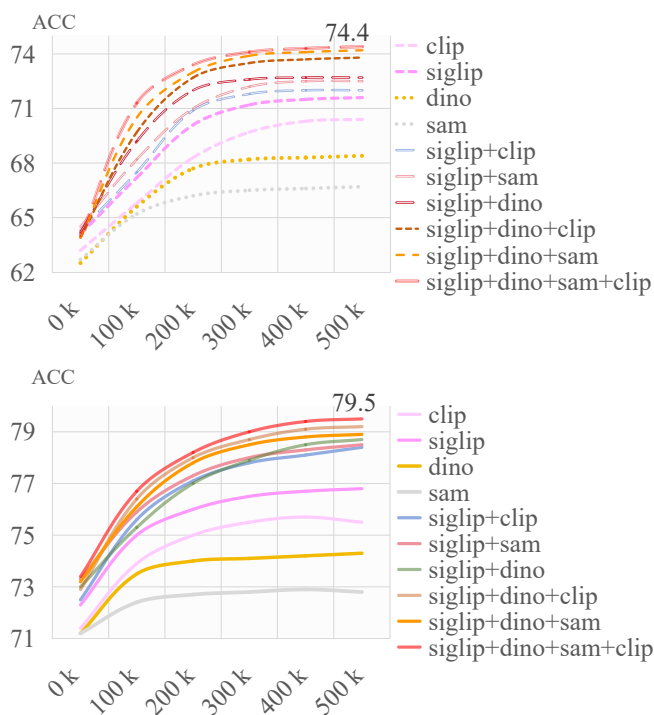


Figure 4: Result of scaling vision model. We post the accuracy of *Test-G* on the top in dashed lines and *Test-S* at the bottom in solid lines.

The 2 “rainbows” in Figure 4 shows the result of experiments on scaling vision model. The top “rainbow” illustrates the performance on *Test-G* set during the accumulation of pre-train data. In the comparison of the four individual visual backbones, SigLIP performed the best. Then CLIP performed the next, followed by DINOv2 and SAM. This reflects the superiority of language-supervised visual backbones in such multimodal tasks. Then we freeze SigLIP and append other visual features. In the comparison of binary

visual backbone combinations, SigLIP+DINO emerged as the fastest and best learner in deep green dashed line. This indicates that the self-supervised backbone DINOv2 helps provide more detailed information, which aids the model in focusing on finer visual details in VSR QA scenarios. We freeze the SigLIP+DINO and add CLIP and SAM respectively. In the comparison of triple backbones, CLIP brought fewer benefits compared to SAM. This may be due to the high feature overlap between CLIP and SigLIP, as they are both contrastive learning models supervised by language. Finally, with the combined efforts of the four backbones, the accuracy of *Test-G* reached the highest value of 74.4.

Similarly, the bottom “rainbow” in Figure 4 shows the accuracy on the *Test-S*. The combination results of each backbone are generally consistent with those of *Test-G*, except that in the comparison of triple backbones, the appended CLIP outperforms SAM. Another point is that compared to *Test-G*, the model learns faster under the single backbone setup. This is likely due to the reduced requirements on the language question side with the use of a fixed template, which lowers the overall task difficulty. Consequently, a smaller amount of data is sufficient for the model to achieve learning saturation. Nevertheless, the scaling model method ultimately improved the accuracy of *Test-S* to 79. Moreover, it can be observed that a common characteristic of both test sets is that while adding new visual features results in accuracy gains, the incremental improvement in accuracy diminishes as the number of added visual features increases. Finally, we designated the best-performing model with 4 backbones trained by the total data obtained as a Visual Spatial Reasoning Expert (VSRE).

Result on Other Benchmarks

To verify the generalization of our expansion method, we tested our VSRE on the MME⁶, MMBench⁷ and SEEDv2⁸. Table 3 show the comparison of VSRE on the related sub-

⁶<https://github.com/BradyFU/Awesome-Multimodal-Large-Language-Models>

⁷<https://mmbench.opencompass.org.cn/home>

⁸<https://github.com/AILab-CVC/SEED-Bench>

Model	MME	MMBench	SEEDv2
MiniGPT-4v2	43.33	-	32.6
Qwen-VL(chat)	128.33	47.2	40.3
LLaVA1.5 13B	133.33	57.6	38.5
BLIP2	73.33	58.4	36.2
VSRE	155.00	64.8	46.6

Table 3: The comparison results of VSRE on the related subsets of other datasets including MME, MMBench and SEEDv2.

sets of other benchmarks. Firstly, on the MME dataset, which also involves answering binary questions as VSR, VSRE achieved the highest score of 155.33, an improvement of more than 22 points compared to the optimized baseline LLaVA1.5 13B. This indicates that our expanded data and model are not overfitted to the single VSR dataset but show robust performance across benchmarks. Moreover, even though MMBench and SEEDv2 ask multiple-choice questions, VSRE still achieved the best results. Compared to the baseline LLaVA before optimization, the scores improved by 7.2 and 8.1 respectively. This may be due to our diverse template-designed IFT dataset, which included question formats similar to multiple-choice questions. Overall, the best performance among the dataset subsets demonstrates that our expansion method for VSR is not overfitting to similar data. Instead, it has genuinely learned visual position reasoning capabilities, handled more diverse text question formats and discerned visual positional information.

More Sensitive Vision Features

To verify that the model has become more sensitive to positional information alongside the accuracy increase, we selected 7 common positional relationships. For each one, we sampled 200 instances and used sklearn’s t-SNE for dimensionality reduction on visual tokens to plot the distribution of each relationship with more details of comparison in the Appendix. It is evident that after dimensionality reduction, the visual features extracted by VSRE are better distinguished, with more pronounced inter-class differences.

Furthermore, Table 4 presents the normalized average intra-class distance for each category. We use this metric to represent the model’s ability to summarize and generalize single positional concepts; a smaller value indicates that the extracted features are better clustered (with more details shown in Appendix). It is evident that the visual features extracted by the more professional spatial sensor VSRE have a smaller average intra-class distance. This indicates that VSRE has a better understanding and generalization capability for various visual spatial concepts.

Bias Result

We re-testing VSRE “yes” and “no” question accuracy with the same template dataset in re-evaluation, and found that the bias issue was alleviated (more details in Appendix). Although the overall accuracy for “yes” questions is still higher than for “no”, the gap between the two has significantly narrowed. As the accuracy of the red “no” questions has signif-

Relations	above	next to	behind	inside
LLaVA(51.2%)	0.54	0.66	0.53	0.42
VSRE(79.5%)	0.24	0.34	0.27	0.21
Relations	on top of	under	on	AVG
LLaVA(51.2%)	0.62	0.57	0.69	0.57
VSRE(79.5%)	0.33	0.29	0.36	0.29

Table 4: Statistic result of average intra-class distance for each spatial relation category on 200 samples by llava1.5 13B (acc 51.2%) and VSRE(acc 79.5%).

icantly improved, VSRE is not misled by the co-occurrence of related entity concepts in text and images but is more focused on the positional relationships between entities.

Conclusion

In this work, we first re-evaluate the VSR ability of VLLMs from scratch and diagnosed issues of inconsistent performance, hypersensitivity on text prompts, lack of perception on visual details and answer bias. To address the problems, we first proposed the unified evaluation test set (*Test-G* and *Test-S*) for VLLM in QA scenario. Next, we proposed methods for expanding both the data and the model structure. We validated the effectiveness of the methods through experiments on scaling data and scaling models. We also validated the generalizability of our method on different benchmarks(MME, MMBench and SEED) and various models(VLLMs and LLMs). Then, our proposed spatial expert VSRE surpasses the performance of LLaVA1.5 13B by 27% (from 52% to 79%) on VSR test set.

Furthermore, in the sensitivity analysis, we found that while VSRE exhibited better differentiation of visual positional concepts, it also demonstrated a superior ability to summarize them. A more sensitive visual position extractor brought about more specialized visual reasoning capabilities. This led the model to focus more on visual positional information, making it less influenced by the co-occurrence of entities and alleviating the issue of biased answer. Ultimately, our expansion method highlights the immense potential in VSR field for related data and model structures.

Acknowledgements

This work was supported by the National Key R&D Program of China under grant 2023YFC3804600 and the Shenzhen Science and Technology Plan - Major Science and Technology Projects (No.KJZD20230923114213027)

References

- Chen, J.; Zhu, D.; Shen, X.; Li, X.; Liu, Z.; Zhang, P.; Krishnamoorthi, R.; Chandra, V.; Xiong, Y.; and Elhoseiny, M. 2023. MiniGPT-v2: large language model as a unified interface for vision-language multi-task learning. arXiv:2310.09478.
- Dai, W.; Li, J.; Li, D.; Tiong, A. M. H.; Zhao, J.; Wang, W.; Li, B.; Fung, P.; and Hoi, S. 2023. InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning. arXiv:2305.06500.

- Fu, C.; Chen, P.; Shen, Y.; Qin, Y.; Zhang, M.; Lin, X.; Yang, J.; Zheng, X.; Li, K.; Sun, X.; Wu, Y.; and Ji, R. 2024. MME: A Comprehensive Evaluation Benchmark for Multimodal Large Language Models. *arXiv:2306.13394*.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. 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*, 6904–6913.
- Karamcheti, S.; Nair, S.; Balakrishna, A.; Liang, P.; Kollar, T.; and Sadigh, D. 2024. Prismatic vlms: Investigating the design space of visually-conditioned language models. *arXiv preprint arXiv:2402.07865*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. *arXiv:2304.02643*.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123: 32–73.
- Li, B.; Wang, R.; Wang, G.; Ge, Y.; Ge, Y.; and Shan, Y. 2023. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, F.; Emerson, G.; and Collier, N. 2023. Visual Spatial Reasoning. *arXiv:2205.00363*.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26296–26306.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024b. Improved Baselines with Visual Instruction Tuning. *arXiv:2310.03744*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023a. Visual Instruction Tuning. *arXiv:2304.08485*.
- Liu, H.; Xue, W.; Chen, Y.; Chen, D.; Zhao, X.; Wang, K.; Hou, L.; Li, R.; and Peng, W. 2024c. A survey on hallucination in large vision-language models. *arXiv preprint arXiv:2402.00253*.
- Liu, Y.; Duan, H.; Zhang, Y.; Li, B.; Zhang, S.; Zhao, W.; Yuan, Y.; Wang, J.; He, C.; Liu, Z.; et al. 2023b. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- McKinzie, B.; Gan, Z.; Fauconnier, J.-P.; Dodge, S.; Zhang, B.; Dufter, P.; Shah, D.; Du, X.; Peng, F.; Weers, F.; et al. 2024. Mm1: Methods, analysis & insights from multimodal llm pre-training. *arXiv preprint arXiv:2403.09611*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; Assran, M.; Ballas, N.; Galuba, W.; Howes, R.; Huang, P.-Y.; Li, S.-W.; Misra, I.; Rabbat, M.; Sharma, V.; Synnaeve, G.; Xu, H.; Jegou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2024. DINOv2: Learning Robust Visual Features without Supervision. *arXiv:2304.07193*.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. *arXiv:2307.01952*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. *arXiv:2103.00020*.
- Tong, S.; Brown, E.; Wu, P.; Woo, S.; Middepogu, M.; Akula, S. C.; Yang, J.; Yang, S.; Iyer, A.; Pan, X.; et al. 2024a. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.
- Tong, S.; Liu, Z.; Zhai, Y.; Ma, Y.; LeCun, Y.; and Xie, S. 2024b. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9568–9578.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid Loss for Language Image Pre-Training. *arXiv:2303.15343*.
- Zhao, X.; Li, X.; Duan, H.; Huang, H.; Li, Y.; Chen, K.; and Yang, H. 2024. MG-LLaVA: Towards Multi-Granularity Visual Instruction Tuning. *arXiv preprint arXiv:2406.17770*.