

COARSE CORRESPONDENCES Boost Spatial-Temporal Reasoning in Multimodal Language Model

Benlin Liu^{1*}, Yuhao Dong^{2,3*}, Yiqin Wang^{2*},

Zixian Ma¹, Yansong Tang², Luming Tang⁵, Yongming Rao³, Wei-Chiu Ma^{5,6}, Ranjay Krishna^{1,5}

¹University of Washington, ²Tsinghua University, ³Tencent,

⁴Allen Institute for AI, ⁵Cornell University

[coarse-correspondence.github.io](https://github.com/coarse-correspondence)

Abstract

Multimodal language models (MLLMs) are increasingly being applied in real-world environments, necessitating their ability to interpret 3D spaces and comprehend temporal dynamics. Current methods often rely on specialized architectural designs or task-specific fine-tuning to achieve this. We introduce COARSE CORRESPONDENCES, a simple lightweight method that enhances MLLMs’ spatial-temporal reasoning with 2D images as input, without modifying the architecture or requiring task-specific fine-tuning. Our method uses a lightweight tracking model to identify primary object correspondences between frames in a video or across different image viewpoints, and then conveys this information to MLLMs through visual prompting. We demonstrate that this simple training-free approach brings substantial gains to GPT4-V/O consistently on four benchmarks that require spatial-temporal reasoning, including +20.5% improvement on ScanQA, +9.7% on OpenEQA’s episodic memory subset, +6.0% on the long-form video benchmark EgoSchema, and +11% on the R2R navigation benchmark. Additionally, we show that COARSE CORRESPONDENCES can also enhance open-source MLLMs’ spatial reasoning (by +6.9% on ScanQA) when applied in both training and inference and that the improvement can generalize to unseen datasets such as SQA3D (+3.1%). Taken together, we show that COARSE CORRESPONDENCES effectively and efficiently boosts models’ performance on downstream tasks requiring spatial-temporal reasoning.

1. Introduction

Intelligence is multi-faceted. While multi-modal language models [28] have shown remarkable linguistic, logical and even mathematical intelligence, many remain doubtful about their visual and spatial intelligence. Despite their excel-

lent performance on visual-linguistic tasks, many recent works [24, 37] demonstrate that state-of-the-art MLLMs still struggle at 3D and long video understanding benchmarks, performing only marginally better than blind text-only baselines. These results suggest that spatial-temporal reasoning is a major bottleneck on MLLMs’ path to general visual intelligence.

To enhance MLLMs’ 3D understanding, researchers have mainly explored three approaches: providing MLLMs with 3D data as input [45], designing specialized architectures for 3D tasks [13], or employing supervised fine-tuning with 3D data [6]. Similarly, to boost MLLMs’ temporal understanding, prior works have proposed new model architectures designed for long video understanding [4, 30], or adopted Socratic-based methods [15, 50] (i.e., converting each frame of a video into text using a caption model, and then using text-only LLMs to summarize).

In contrast to prior works that separately enhance spatial and temporal reasoning, we propose COARSE CORRESPONDENCES, a simple yet effective training-free visual prompting method to jointly boost spatial-temporal reasoning in MLLMs. COARSE CORRESPONDENCES uses a tracking model to extract object-level correspondences across multiple images, and then represent the most salient correspondence relationships on the images through visual prompting. We refer to our method as *Coarse* because of the following:

1. We only visually prompt for instance-level correspondences and not point-level correspondences.
2. The instance-level correspondences are extracted using off-the-shelf tracking models.
3. We only visualize a handful of prominent corresponding instances.

Despite the obtained correspondence information not being perfectly precise, our method can still significantly boost MLLMs’ spatial-temporal reasoning with only 2D image inputs and without any specialized architectural design or task-specific fine-tuning.

*Equal contribution.

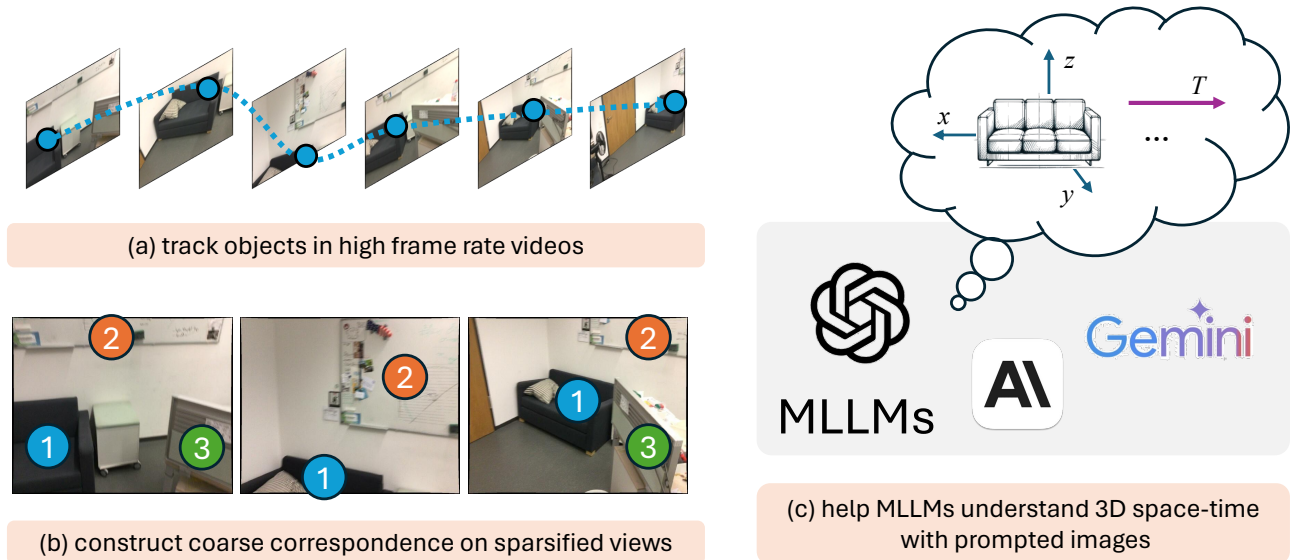


Figure 1. We combined light-weight video tracking models and multimodal LLMs to achieve a better understanding of 3D spacetime. (a) We use a tracking model at a high frame rate to obtain instance segmentation masks for each frame. (b) Then, we sequentially sparsify input frames, select prominent coarse correspondences, and visualize the constructed coarse correspondences on the images. (c) Finally, we enable MLLMs to better understand 3D spacetime from the prompted images.

We have demonstrated substantial performance gains of COARSE CORRESPONDENCES through extensive experiments with both open-source and closed-source models across 6 benchmarks on spatial-temporal reasoning. For closed-source models, we apply COARSE CORRESPONDENCES on GPT4-V/O during inference and achieve compelling gains. First, on spatial reasoning, we show that our method significantly surpasses state-of-the-art models by 20.5% and 9.7% on ScanQA [2] and OpenEQA [24] respectively. Second, for long video understanding, our method leads to a 6% gain in performance on the EgoSchema benchmark [25]. Notably, our method uses much fewer input images and, in a zero-shot manner, outperforms many fine-tuned models that use far more images. For example, on EgoSchema, COARSE CORRESPONDENCES surpasses state-of-the-art results with just 8 uniformly sampled frames from a 3-minute video, greatly reducing the computational costs of MLLMs compared to existing methods. In addition to 3D and video QA tasks, we further demonstrate that our method enhances models’ performance on embodied tasks such as navigation [16], which require strong spatial-temporal reasoning, by 11% in success rate on R2R. These results suggest that COARSE CORRESPONDENCES boosts MLLMs’ spatial-temporal reasoning both effectively and efficiently. Last but not least, we experiment with open-source MLLMs [22] by applying COARSE CORRESPONDENCES in both instruction tuning and inference; again, our method shows performance gains against the baseline (by 6.9% on ScanQA), and the improvement even generalizes to unseen datasets such as SQA3D (+3.1%). These results suggest that COARSE COR-

RESPONDENCES works well universally with any model – both closed-source and open-source – that can take in multiple images and understand visual markers.

Overall, we want to highlight with this work that, despite its simplicity and being underestimated for semantic tasks in deep learning, visual correspondence can still bring significant utility to spatial-temporal reasoning in MLLMs, just as it has long contributed to 3D reconstruction [35]. We hope our work demonstrates the potential of leveraging general-purpose MLLMs to better understand our physical world.

2. Method

We introduce COARSE CORRESPONDENCES, a visual prompting method that allows MLLMs to reason about 3D space and time.

Problem formulation. Given a question Q and a sequence or set of observations in an environment $[I_1, \dots, I_n]$, our aim is to design a visual prompt $\mathcal{P}(\dots)$ that modifies the input image set. These image inputs don’t have to be a video. They can also represent a set of images of a specific scene from multiple viewpoints. We evaluate the prompt by measuring its utility in prompting an MLLM \mathcal{M} :

$$[I'_1, \dots, I'_n] = \mathcal{P}([I_1, \dots, I_n])$$

$$\hat{\mathcal{A}} = \mathcal{M}([I'_1, \dots, I'_n], Q)$$

We compare generated answer $\hat{\mathcal{A}}$ with the ground truth \mathcal{A} .

In our framework, the MLLMs can be any general-purpose model without requiring any special architecture

or training for spatial-temporal reasoning. Our aim is to develop a prompting strategy that allows models to improve such capabilities without any training (Figure 1).

Coarse Correspondence

Our prompting method, COARSE CORRESPONDENCES, contains four steps: (1) tracking correspondences, (2) sparsify frames, (3) selecting, and (4) visualizing coarse correspondences.

(1) Tracking correspondences. Given n input images, $[I_1, \dots, I_n]$, we first use an off-the-shelf video object tracking model, such as Tracking Anything [47]. This model extracts class-agnostic instance segmentation masks (M_1, \dots, M_n) for each image. Each M_i is a $H \times W$ dimensional matrix where H and W are the height and width of the input image I_i . Each pixel location in M_i contains an instance ID, indicating which instance the pixel at that position belongs to within the image sequence.

(2) Sparsify frames. Since most MLLMs contain a large number of parameters, directly using them to process long image sequences is very computationally intensive. Additionally, proprietary MLLMs like GPT-4O can also incur significant costs if the number of image tokens that need to be processed increases. Reducing the number of input images might lose vital information necessary for MLLMs.

COARSE CORRESPONDENCES strikes a balance in this tradeoff by extracting meaningful video object tracks (a relatively cheaper operation) from high-frame-rate image sequences, and then samples a few image inputs along with the tracks, to retain—and even improve—performance while reducing the MLLM’s computation cost. From this extracted video object tracks, we perform temporal downsampling, retaining only $m \ll n$ uniformly sampled images and their corresponding masks, denoted as $[I_{s_1}, \dots, I_{s_m}]$ and $[M_{s_1}, \dots, M_{s_m}]$, where $s_i \in \{1, \dots, n\}$. This downsampling reduces the number of images we feed into \mathcal{M} .

(3) Selecting coarse correspondences. Prompting an MLLM with all the detected correspondences results in information overload. In fact, our ablations (discussed in Sec 5) find that adding all the correspondences reduces the MLLM’s performance. Therefore, we select a subset of prominent instances to retain. We select the prominent instances of the top- K objects that co-occur in the most number of frames. We first calculate the occurrence frequency and area sum of each unique instance ID in the retained m masks using the following equation:

$$Freq(\text{ID}) = \sum_{i=s_1}^{s_m} \mathbf{1}_{\{\text{ID} \in M_i\}}, Area(\text{ID}) = \sum_{i=s_1}^{s_m} \sum_{p \in M_i} \mathbf{1}_{\{\text{ID}=p\}}.$$

Then, we first sort all instance IDs in descending order based on $Freq(\text{ID})$. If there are ties, we further sort based on $Area(\text{ID})$. Finally, we retain the top k instance IDs as tracklets, denoted as $[T_1, \dots, T_k]$, to visualize for MLLMs.

(4) Visualizing coarse correspondences. For each set of obtained correspondence relationships, we visualize the correspondences directly in the image as a marker. Specifically, for each identified primary instance ID T_i , if it exists in the mask M_{s_j} of a retained image I_{s_j} , we overlay a mark with a fixed size and shape labeled with T_i at the position $(\bar{x}_{ij}, \bar{y}_{ij})$ on I_{s_j} to produce I'_{s_j} . The specific placement position can be easily obtained by:

$$(\bar{x}_{ij}, \bar{y}_{ij}) = \frac{\sum_{(x,y)} (x, y) \cdot \mathbf{1}_{\{M_{s_j}(x,y)=T_i\}}}{\sum_{(x,y)} \mathbf{1}_{\{M_{s_j}(x,y)=T_i\}}}$$

Naturally, we can overlay not just the markers but also the segmentation outlines or even the segmentation masks associated with each retained prominent instance. We explore these ablations later. In the end, we obtain the prompted image sequence $[I'_1, \dots, I'_m]$, which is then used as the input to MLLMs.

3. Prompting Proprietary Models

We first evaluated the utility of our COARSE CORRESPONDENCES on multiple tasks using proprietary models, including understanding 3D space (ScanQA [2] and OpenEQA [24] in §3.1) as well as temporal events (EgoSchema [25] in §3.2). Building on the improvements our method brings to spatial-temporal reasoning, we further demonstrate that our method also delivers significant gains in navigation tasks (VLN-CE [16]). Across all these benchmarks, we augment proprietary MLLMs (e.g., GPT-4V and GPT-4O) with COARSE CORRESPONDENCES and evaluate its zero-shot performance. We show that COARSE CORRESPONDENCES significantly improves the base GPT models and can substantially surpass many current state-of-the-art methods that have undergone specialized fine-tuning, even while using much fewer images as input. All experiments were conducted using A100 80G GPUs.

3.1. Spatial understanding

Benchmarks. The validation set of ScanQA dataset contains 4675 questions about 71 scenes. The questions require basic recognition, 3D localization, and 3D embodied capabilities [9]. It contains two ground-truth answers per question for evaluation with models that produce free-form answers. OpenEQA Dataset is an open-vocabulary dataset benchmarking spatial environment understanding and embodied reasoning. We evaluate on OpenEQA’s EM-EQA data split, which contains over 1600 high-quality human-generated questions. The subset tests the episodic memory of an agent moving through a 3D environment over time.

Baselines. For ScanQA, we evaluate COARSE CORRESPONDENCES by augmenting both GPT-4{V,O}, Gemini and Claude models. Besides, we also consider 3D specialized models [2, 13, 49] fine-tuned on ScanQA. For OpenEQA,

| Model | Frame | BLEU-1 | BLEU-2 | METEOR | ROUGE-L | CIDEr |
|---------------------------------------|-------|-------------|-------------|-------------|-------------|-------------|
| <i>3D-Specific Models</i> | | | | | | |
| ScanQA [2] | - | 26.9 | 16.6 | 11.5 | 30 | 55.4 |
| ScanRefer+MCAN [49] | - | 30.2 | 20.4 | 13.1 | 33.3 | 64.9 |
| 3D-LLM [13] | - | 39.3 | 25.2 | 14.5 | 35.7 | 69.4 |
| <i>Open-source Multi-modal Models</i> | | | | | | |
| LLaVA(Fine-tuned) | 64 | 34.7 | 22.0 | 13.8 | 31.1 | 67.3 |
| LLaVA+Coarse Correspondences | 64 | 38.6 | 24.7 | 15.4 | 38.3 | 74.2 |
| <i>Proprietary Multi-modal Models</i> | | | | | | |
| GPT-4V | 8 | 28.6 | 13.4 | 13.5 | 33.4 | 59.6 |
| GPT-4V+Coarse Correspondences | 8 | 39.7 | 25.5 | 17.4 | 40.8 | 79.2 |
| GPT-4O | 4 | 30.5 | 19.8 | 14.8 | 36.1 | 72.2 |
| GPT-4O+Coarse Correspondences | 4 | 35.4 | 25.5 | 18.0 | 42.6 | 87.0 |

Table 1. **Comparison on ScanQA validation set.** We conduct experiments on the ScanQA validation set to demonstrate the effectiveness of COARSE CORRESPONDENCES with different MLLMs. Our method enables both proprietary models and open-source models to surpass all 3D-specific models.

we compare against language-only models to account for language bias (LLaMA2 [40]), commonly used general-purpose multimodal LLMs (GPT-4 [29], Claude3 [1], Gemini-Pro [38], GPT-4V with 15 and 50 frames).

Metrics. For ScanQA, following prior works, we adopt BLEU [31], METEOR [5], ROUGE-L [20], and CIDEr [41] as our evaluation metrics. For OpenEQA, we follow their evaluation approach by using GPT-4 to compare the generated answers with the ground-truth answers and assign a score. We report the average score across all questions.

Results. For ScanQA, as shown in Table 1, compared to raw input, COARSE CORRESPONDENCES consistently improves the overall performance of different proprietary models. For instance, on the strongest model, GPT-4o, COARSE CORRESPONDENCES brings improvements of 5.7 BLEU-2, 3.2 METEOR, 6.5 ROUGE-L, and 15 CIDEr points. Compared to methods that are specifically designed for 3D understanding tasks, fine-tuned with specialized 3D SFT data, or even those that use 3D point clouds instead of 2D images as input, we observe that a general-purpose MLLM can still outperform them, especially when enhanced with COARSE CORRESPONDENCES. Moreover, we found that this can be achieved using far fewer images as input.

We also demonstrated the same conclusion on OpenEQA, as indicated in Table 2. By applying COARSE CORRESPONDENCES, we significantly improved the performance of both GPT-4v and GPT-4o, achieving better results with fewer input images. These findings suggest that general-purpose MLLMs are indeed capable of understanding 3D space, and COARSE CORRESPONDENCES can significantly enhance their spatial understanding while simultaneously reducing the number of views needed, which could lower the inference cost and make MLLMs more useful for embodied tasks.

| Models | Frame | Accuracy |
|------------------|-------|-------------|
| LLaMA2 [40] | 0 | 28.3 |
| GPT-4 [29] | 0 | 33.5 |
| Claude3 [1] | 20 | 36.3 |
| Gemini-Pro [38] | 15 | 44.9 |
| GPT-4V [27] | 15 | 54.6 |
| GPT-4V [27] | 50 | 55.3 |
| Human | Full | 86.8 |
| GPT-4V | 8 | 44.8 |
| GPT-4V+CC | 8 | 58.5 |
| GPT-4O | 4 | 49.4 |
| GPT-4O+CC | 4 | 59.1 |

Table 2. **Comparisons on EM-EQA setting of OpenEQA.** Our method further enhances the embodied ability of MLLMs and exceeds previous methods by a large margin.

3.2. Temporal understanding

Benchmarks. We evaluated the improvements of our method for long video understanding using the EgoSchema dataset. Each video in EgoSchema is 3 minutes long, with a corresponding question that includes five multiple-choice options. These questions are designed to ensure that answering them requires viewing a sufficient number of frames from the video. Due to budget constraints, we limited our evaluation to 500 questions from the validation set.

Baselines. The baseline methods we compared against include newly designed and trained model architectures specifically for long video understanding, such as LongViviT [30] and MC-ViT-L [4]. On the other hand, we also compared methods that rely on text-only foundation models (e.g., GPT-4), i.e., Socratic-based approaches [15, 50], which first use an

| Models | Frame | Accuracy |
|-----------------|-------|-------------|
| LongViviT [30] | 256 | 56.8 |
| MC-ViT-L [4] | 128+ | 62.6 |
| LLoVi [50] | 180 | 58.3 |
| VideoAgent [42] | 8.4 | 60.2 |
| MVU [33] | 16 | 60.3 |
| VideoAgent [10] | - | 62.8 |
| LangRepo [15] | - | 66.2 |
| GPT-4V | 8 | 64.2 |
| GPT-4V+CC | 8 | 67.4 |
| GPT-4O | 8 | 67.2 |
| GPT-4O+CC | 8 | 73.2 |

Table 3. **Comparisons on EgoSchema validation set.** COARSE CORRESPONDENCES improves existing MLLMs and surpasses previous fine-tuned models in a zero-shot manner.

off-the-shelf image captioning model [51] to convert video frames into captions, and then prompt GPT-4 to answer questions based on those captions. Additionally, we compared agent-based methods [10, 42], which involve using GPT-4 alongside an image captioning model in an agent framework to iteratively perform a series of multi-step reasoning operations to understand long videos. In contrast to these approaches, our method is entirely based on an end-to-end general MLLM architecture, exploring how to further enhance its ability to understand long videos.

Results. COARSE CORRESPONDENCES demonstrates state-of-the-art performance, significantly outperforming existing approaches in a zero-shot manner (Table 3). Compared to the original GPT-4o model, our method improves its performance by 6%. Notably, our method uses far fewer frames than other approaches, yet achieves higher results compared to methods that use many more frames. It is also worth highlighting that even the original GPT-4o, when limited to just 8 frames, already serves as a very strong baseline. This indicates the potential of a sufficiently powerful general-purpose MLLM in long video understanding. We provide a detailed breakdown of COARSE CORRESPONDENCES’s performance across different question types in the supplementary material.

3.3. Navigation

Building on the improvements in 3DQA and VideoQA, we hope that our method can also prove effective in embodied tasks such as navigation. Navigation requires an agent to understand 3D space, such as being able to determine the spatial relationship between objects in the instruction and itself, while also performing temporal reasoning to assess the progress toward completing the instruction. We consider conducting experiments on the VLN-CE benchmark [16], which is a continuous simulation environment for low-level action execution in indoor scenes.

Setup. We adopt the val-unseen split from R2R [16] for evaluation. Unlike the previous QA tasks, where all images could be processed at once, in navigation tasks, each image is processed in an online fashion. Specifically, we feed in one image at each iteration of the conversation. Given the significant variation in viewpoints during navigation, we use SAMv2 [34], the state-of-the-art model for long-range tracking, to label each new input image based on episodic history. Then, using the prompted images, we induce the MLLM to output one of four actions at each step: FORWARD (distance), TURN-LEFT (rotation angle), TURN-RIGHT (rotation angle), and STOP. We follow NavGPT [52] to craft input prompts. Considering the high computational cost of navigation tasks, we selected 100 samples from the val-unseen split. Our primary goal is to demonstrate that our method can enhance GPT models’ capabilities in zero-shot navigation tasks, which remains a significant challenge for various types of models.

Metrics. We follow the standard VLN evaluation metrics to evaluate the navigation performance, including success rate (SR), oracle success rate (OS), success weighted by path length (SPL), trajectory length (TL), and navigation error from goal (NE). Note that an episode is considered successful if the agent calls the STOP action within 3m of the goal in the VLN-CE.

Results. As shown in Table 4, our method achieved improvements across all metrics. While GPT-4o performs impressively on many QA tasks, its zero-shot performance on navigation tasks is relatively low. This may partly be due to the lack of specialized training on action data, making it less accurate in outputs such as determining how many meters to move forward. However, our method reveals another dimension of the problem: MLLMs’ understanding of the 3D spacetime in which they operate can be further enhanced. This is evidenced by the significant improvements in navigation when using COARSE CORRESPONDENCES. We believe that our approach holds great potential for embodied tasks, which can be explored in future research.

4. Prompting Open Models

We further validate the effectiveness of our COARSE CORRESPONDENCES on open-source models. Our primary goal is to demonstrate that our method not only enhances proprietary models but also yields general improvements across various MLLMs. However, directly prompting open-source models presents two challenges: first, these models lack the ability to interpret prompting marks as effectively as proprietary models; second, they currently cannot handle multiple image inputs well. Thus, we first fine-tune the open-source models to understand prompting marks and manage multiple image inputs, enabling COARSE CORRESPONDENCES to enhance these models’ environmental understanding capabilities at the inference stage. In this work, we primarily focus

| Methods | Success Rate \uparrow | Oracle Success Rate \uparrow | Success weighted by Path Length \uparrow | Trajectory Length \uparrow | Navigation Error \downarrow |
|-----------|-------------------------|--------------------------------|--|------------------------------|-------------------------------|
| GPT-4O | 12.00 | 18.00 | 10.37 | 7.31 | 8.49 |
| GPT-4O+CC | 23.00 | 29.00 | 21.03 | 8.12 | 7.37 |

Table 4. **Comparison on Navigation task.** COARSE CORRESPONDENCES improve GPT-4o’s performance on R2R dataset for different evaluation metrics.

| Method | Acc |
|-------------------------------|-------|
| LLaVA(w/o CC during finetune) | 36.0 |
| LLaVA+CC(train only) | 37.17 |
| LLaVA+CC(train & test) | 39.13 |

Table 5. **Comparisons on SQA3D dataset.** We tested the effectiveness of COARSE CORRESPONDENCES on the out-of-domain SQA3D dataset. With the same training data, we found that using COARSE CORRESPONDENCES not only improves performance during inference but also enhances zero-shot evaluation when applied solely during the training phase.

on proving that COARSE CORRESPONDENCES improves open-source models’ understanding of 3D space.

We start with the LLaVA model [22] and fine-tune it using a dataset comprising approximately 1.2 million samples of image and video data. Notably, the ScanQA dataset is the *only* 3D-related dataset in this collection. COARSE CORRESPONDENCES is applied exclusively to the ScanQA data, while other data maintain their original format. Fine-tuning on such a dataset enables the open-source LLaVA model to better interpret overlaid marks on images and process multiple image inputs.

In-domain Evaluation. We first evaluate whether COARSE CORRESPONDENCES can improve open-source models on the ScanQA validation set. As shown in Table 1, our method significantly enhances the model’s 3D spatial understanding compared to the baseline fine-tuned on the same data without using COARSE CORRESPONDENCES, even outperforming previous VLMs [13] specifically designed for 3D tasks, which incorporate specialized architectural designs and are fine-tuned on substantially larger amounts of 3D-related data. This in-domain evaluation demonstrates that COARSE CORRESPONDENCES is also effective for open MLLMs.

Out-domain Evaluation. To further demonstrate the generalizability of our method, we conduct experiments on the SQA3D dataset to evaluate whether COARSE CORRESPONDENCES can also provide improvements in a zero-shot setting. As shown in Table 5, on this previously unseen dataset, COARSE CORRESPONDENCES still outperforms the baseline simply fine-tuned on the same data, proving that our method can enhance open-source models’ understanding of 3D space on out-of-domain datasets.

Notably, even without using COARSE CORRESPON-

DENCES during inference, simply applying it during training already yields improvements. This highlights that our method is effective not only as a prompting technique for inference but also holds potential as a data augmentation method during training, a prospect worth exploring further in future work.

5. Analysis

Here, we explore the various design decisions in our method. Although the study is conducted only on ScanQA, applying the design choices derived from ScanQA to all other benchmarks also yields significant performance improvements. This strongly demonstrates the generalizability of these design choices. Additionally, we further analyze the impact of our method on mitigating the camera motion bias present in MLLMs.

How does COARSE CORRESPONDENCES differ from other visual prompting methods? Our proposed method calculates and highlights correspondences between images, aiming to elicit 3D and temporal understanding. Other visual prompting methods (namely Set-of-Mark [48], 3DaxesPrompts [21], and Chain-of-thought [43]) can also be viewed as alternative prompting methods. Given that the ground-truth answers in existing benchmarks are relatively brief, we selected a scene from ScanQA and manually designed a new question. We qualitatively compare COARSE CORRESPONDENCES against other prompting methods on this new question, as shown in Figure 2.

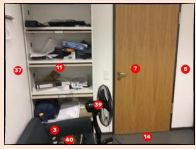
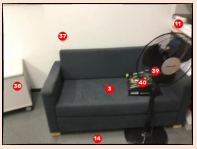
The orange part of Figure 2 shows our **Coarse Correspondence** labels are recognized by GPT-4V. The output answer provides evidence that our coarse correspondence helps GPT-4V develop a mental 3D model of the scene. **Set-of-Marks** provides no spatial corresponding information and therefore is unhelpful. The Axis labels in **3Daxes** can be easily misrecognized by GPT-4V, leading to misleading spatial information. Though Chain-of-Thought helps identify objects, it fails to resolve the “spatial perspective-taking” issue.

Why use coarse instead of dense correspondences? Instead of filtering and retaining only a handful of coarse correspondences, one ablation we considered is the possibility of using all dense correspondence. Unfortunately, we find that excessively overlaying too many instance marks can degrade performance (Figure 3) as they occlude the visual content in the images.

System: You are an AI with the ability to analyze a series of images, each representing a different perspective of a single scene. [Prompt-about-Marks]. Your task is to construct a 3D understanding based on these images.

User: You are sitting on the sofa and the electric fan is on your left. Describe the location of the room door from your perspective: A. to the front left of you; B. to the front right of you; C. to the back left of you; D. to the back right of you.


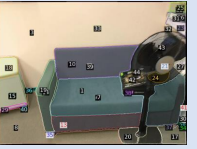
+Coarse Correspondence (Ours)

View 1 View 2

GPT-4V: The first image ... a door marked with number 7, a fan marked with number 39, and items on shelves marked with numbers 11 and 37. In the second image, ... the fan again at number 39, and a sofa marked with number 3(✓).
 We can infer that the door and the fan are on the same side of the room(✓). Since the fan appears to the left in the first image and is also on the left when sitting on the sofa as per the second image and your description, the door must be to the front of the sofa, as it is not visible in the second image from sofa's perspective(✓). By 'sitting on the sofa' as instructed, and with the electric fan to the left, I can determine the door is A. to the front left of you (✓).


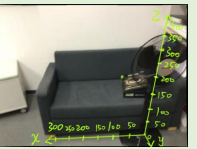
+Set-of-Mark

View 1 View 2

GPT-4V: The fan is seen on the right side from the sofa's perspective, its leftward when sitting means the orientation is flipped(✓). Therefore, the door that is visible on the right in the first image would be D. behind and to the right when seated on the sofa(X).

+3DAxesPrompts

View 1 View 2

GPT-4V: The first image shows ... and a coordinate system with X-axis (horizontally rightwards)(X) and Z-axis (vertically upwards). The second image shows a sofa on the left with a coordinate system's Y-axis (horizontally extending away from the viewer) (X) and Z-axis as before. When sitting on the sofa with the fan to your left, the door would be B. diagonally opposite to you on the right side (X).

+Chain-of-Thought

System: Your response should be as follows.
1)Identification: Identify all labeled objects.
2)Analysis: Analyze the objects and their positions construct a 3D model.
3)Synthesis: Combine the info above, form a comprehensive understanding.
4)Response: Provide a concise answer to the question based on your comprehensive 3D modelling.

GPT-4V:
 Identification: Door, shelves, sofa, ... (✓)
 Analysis: in image 1 the door is on the right with the shelves adjacent to the door ... In image 2 the sofa is ... (✓).
 Synthesis: When sitting on the sofa, the shelves are on the right(X). The door is next to the shelves.
 Response: B.(X).

Figure 2. Comparison of different prompting method. Our proposed Coarse Correspondences successfully guided GPT-4O to understand 3D spatial relationships and generate the right answer. Other existing prompting method including image-based Set-of-Marks, 3DAxes and text-based Chain-of-Thought failed to answer correctly.

How large should the marks be? We inject the correspondences into MLLMs by overlaying the marks into images. We empirically find an optimal mark size (where ‘px’ represents the mark’s diameter in pixels) in Figure 3. Marks that are too small tend to be ignored while those that are too large occlude visual content.

What shape should the marks be? We further studied the appearance of the marks. In addition to red circles with white text, we experimented with adding segmentation outlines and segmentation masks. As shown in Figure 3, using segmentation outlines enhances object grounding. However, using segmentation marks occludes visual content and reduces performance.

Improvement of Camera Motion Invariance with COARSE CORRESPONDENCES. A crucial aspect of strong spatial reasoning is the model’s ability to maintain an understanding of spatial relationships invariant to camera motion; in other words, whether the image sequence is captured from left-to-right or right-to-left, the model’s comprehension of static space should remain consistent. However, current benchmarks lack annotations regarding whether the 3D space scan was conducted from left-to-right or right-to-left, making it challenging to analyze the effect of camera motion on MLLMs’ spatial understanding. To address this,

we curated a small diagnostic benchmark specifically for this analysis. This benchmark includes 10 scenes captured from left-to-right, with five questions per scene focusing on left-right spatial relationships between objects—fundamental aspects of spatial understanding. We ensured that all questions required information across multiple frames and could not be answered from a single frame alone. Ideally, regardless of whether the image sequence is input in the original or reversed order, the output should remain consistent. We leave additional details on data curation in the supplementary material.

To mitigate randomness, each question was tested 20 times, yielding a total of 1000 trials, and we report the average accuracy in Table 6. We then reversed the sequence order to simulate a right-to-left capture and repeated 1000 trials with this reversed input, reporting the average accuracy for this reverse sequence. Finally, we report the harmonic mean of the average accuracies for the forward and reverse inputs.

Without COARSE CORRESPONDENCES, the accuracy in understanding spatial relationships in right-to-left sequences was significantly lower than in left-to-right sequences. With COARSE CORRESPONDENCES, not only did MLLMs achieve a substantial increase in spatial un-

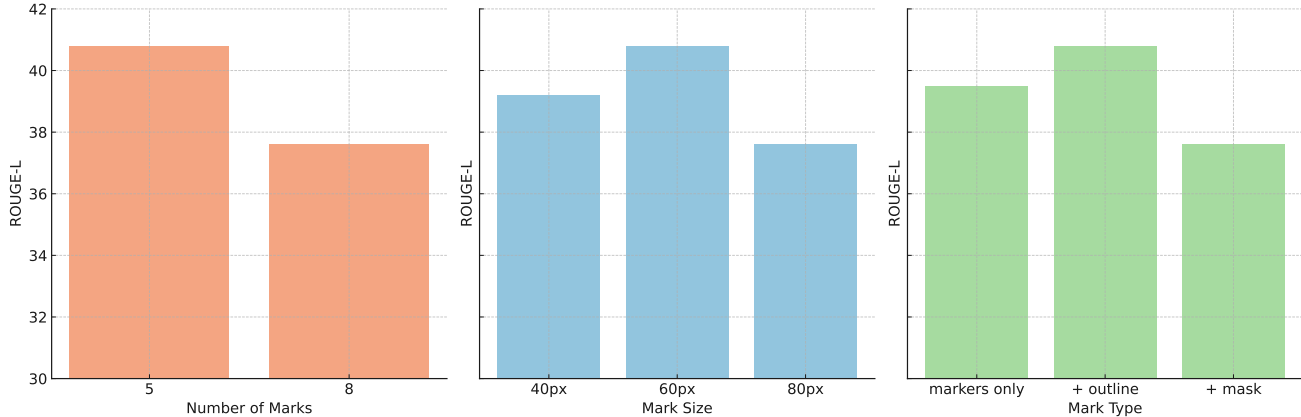


Figure 3. **Ablations on different design choices of COARSE CORRESPONDENCES.** We studied the impact of the number, size, and type of marks on performance. All experiments were conducted on ScanQA using GPT-4V.

| Models | Frame | Forward | Reverse | Harmonic Mean |
|-----------|-------|-------------|-------------|---------------|
| GPT-4O | 4 | 58.0 | 50.4 | 53.9 |
| GPT-4O+CC | 4 | 71.2 | 71.2 | 71.2 |

Table 6. **Effect of COARSE CORRESPONDENCES on Camera Motion Bias.** COARSE CORRESPONDENCES shows strong capability of enhancing 3D spatial understanding of MLLMs. It can also ease the camera motion bias of current MLLMs.

derstanding accuracy with the original sequence (+13.2%), but more importantly, reversing the input order no longer drastically impacted the model’s comprehension of static spatial relationships. Specifically, COARSE CORRESPONDENCES improved the harmonic mean of forward and reverse sequence accuracy by 17.3%. This demonstrates that our method enables MLLMs to perform spatial reasoning invariant to camera motion, thus achieving more robust spatial understanding capabilities.

6. Related work

Multimodal language models Multimodal LLMs [3, 23] integrate vision encoders [32] into large LLMs [8, 40], enabling direct reasoning over visual input. Proprietary models like GPT-4 [27], Gemini [38], and Claude [1], along with open-source models such as the LLaVA series [23] and BLIP series [18], have made significant strides in 2D vision-language tasks like image captioning [7] and visual question answering (VQA) [11, 14]. However, for applications in real-world scenarios like autonomous driving [39] and robotics [46], MLLMs still require stronger spatial-temporal reasoning capabilities. Recent work [19] has enhanced MLLMs’ capacity to process multiple image frames and understand temporal dynamics, achieving notable improvements in video dense captioning [17] and videoQA [12, 44]. Nonetheless, MLLMs still face challenges in understanding complex 3D spatial relations [24] and long-term temporal dynamics [25] present in image sequences. Many

approaches have introduced specialized architectures [13] or fine-tuning methods targeting spatial and temporal reasoning individually. Our goal, however, is to jointly enhance spatial-temporal reasoning within the existing general-purpose MLLM framework. By leveraging off-the-shelf tracking models to extract instance-level correspondences, we achieve this in a training-free manner.

Visual prompting. Effective prompting has been widely proven to improve LLMs across multiple domains. Methods, such as chain-of-thought prompting [43], force the model to reason before answering a question. In the context of visual prompting, we can enhance MLLMs’ grounding abilities [36, 48], their ability to use numerical expressions to describe 3D object relationships [21], and their capability to generate outputs for robotics control [26], all using single-image inputs. Unlike these approaches, COARSE CORRESPONDENCES is designed to enhance MLLMs’ spatial-temporal reasoning capabilities for interpreting complex spacetime structures represented in image sequences.

7. Conclusion.

We propose a framework called COARSE CORRESPONDENCES prompting. By using off-the-shelf video tracking models to obtain class-agnostic, instance-level correspondences and conveying this information to MLLMs through visual prompting, we discovered that this simple method, using only 2D images as input—without any specialized architectural design or task-specific SFT—can effectively enhance MLLMs’ spatial-temporal reasoning. This improvement extends to embodied tasks like navigation. Our method not only works on proprietary models but also generalizes to open-source models, and it performs well on both in-domain and out-of-domain datasets. Moreover, it enhances not just inference but also training. Further analysis shows that our method helps MLLMs become more robust to camera motion bias.

Acknowledgement

We thank Xiaojuan Wang, Ruitao Zhang, Yuntian Deng for helpful discussions, feedback and collecting data. This project is partially funded by Amazon Science. This paper was partially supported by a gift from Ai2 and DARPA Transfer from Imprecise and Abstract Models to Autonomous Technologies (TIAMAT) program No. HR00112490422. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of DARPA.

References

- [1] AI Anthropic. The claude 3 model family: Opus, sonnet, haiku. *Claude-3 Model Card*, 2024. 4, 8
- [2] Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. Scanqa: 3d question answering for spatial scene understanding. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19129–19139, 2022. 2, 3, 4
- [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, 2023. 8
- [4] Ivana Balažević, Yuge Shi, Pinelopi Papalampidi, Rahma Chaabouni, Skanda Koppula, and Olivier J. Hénaff. Memory consolidation enables long-context video understanding, 2024. 1, 4, 5
- [5] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. 4
- [6] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brian Ichter, Danny Driess, Pete Florence, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. *arXiv preprint arXiv:2401.12168*, 2024. 1
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server, 2015. 8
- [8] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, 2023. 8
- [9] Jiafei Duan, Samson Yu, Hui Li Tan, Hongyuan Zhu, and Cheston Tan. A survey of embodied ai: From simulators to research tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 6(2):230–244, 2022. 3
- [10] Yue Fan, Xiaojian Ma, Rujie Wu, Yuntao Du, Jiaqi Li, Zhi Gao, and Qing Li. Videoagent: A memory-augmented multimodal agent for video understanding. *arXiv preprint arXiv:2403.11481*, 2024. 5
- [11] 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, 2017. 8
- [12] Madeleine Grunde-McLaughlin, Ranjay Krishna, and Maneesh Agrawala. Agqa: A benchmark for compositional spatio-temporal reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11287–11297, 2021. 8
- [13] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 3, 4, 6, 8
- [14] Drew A. Hudson and Christopher D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering, 2019. 8
- [15] Kumara Kahatapitiya, Kanchana Ranasinghe, Jongwoo Park, and Michael S. Ryoo. Language repository for long video understanding, 2024. 1, 4, 5
- [16] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*, pages 104–120. Springer, 2020. 2, 3, 5
- [17] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos, 2017. 8
- [18] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 8
- [19] Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023. 8
- [20] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 4
- [21] Dingning Liu, Xiaomeng Dong, Renrui Zhang, Xu Luo, Peng Gao, Xiaoshui Huang, Yongshun Gong, and Zhihui Wang. 3daxiesprompts: Unleashing the 3d spatial task capabilities of gpt-4v. *arXiv preprint arXiv:2312.09738*, 2023. 6, 8
- [22] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *NeurIPS*, 36, 2024. 2, 6
- [23] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 8
- [24] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Sasha Sax, and Aravind Rajeswaran. Openeqa: Embodied question answering in the era of foundation models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2, 3, 8

- [25] Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. Egoschema: A diagnostic benchmark for very long-form video language understanding, 2023. 2, 3, 8
- [26] Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, et al. Pivot: Iterative visual prompting elicits actionable knowledge for vlms. *arXiv preprint arXiv:2402.07872*, 2024. 8
- [27] OpenAI. Gpt-4v(ision) system card. *OpenAI Blog*, 2023. 4, 8
- [28] OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. Accessed: 2024-05-22. 1
- [29] OpenAI, Josh Achiam, and et al. Gpt-4 technical report, 2024. 4
- [30] Pinelopi Papalampidi, Skanda Koppula, Shreya Pathak, Justin Chiu, Joe Heyward, Viorica Patraucean, Jiajun Shen, Antoine Miech, Andrew Zisserman, and Aida Nematzdeh. A simple recipe for contrastively pre-training video-first encoders beyond 16 frames, 2023. 1, 4, 5
- [31] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002. 4
- [32] 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. 8
- [33] Kanchana Ranasinghe, Xiang Li, Kumara Kahatapitiya, and Michael S. Ryoo. Understanding long videos in one multimodal language model pass, 2024. 5
- [34] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 5
- [35] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 2
- [36] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. What does clip know about a red circle? visual prompt engineering for vlms. *arXiv preprint arXiv:2304.06712*, 2023. 8
- [37] Simranjit Singh, Georgios Pavlakos, and Dimitrios Stamoulis. Evaluating zero-shot gpt-4v performance on 3d visual question answering benchmarks. *arXiv preprint arXiv:2405.18831*, 2024. 1
- [38] Gemini Team, Rohan Anil, and et al. Gemini: A family of highly capable multimodal models, 2024. 4, 8
- [39] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Chenxu Hu, Yang Wang, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. Drivevlm: The convergence of autonomous driving and large vision-language models, 2024. 8
- [40] 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. 4, 8
- [41] 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. 4
- [42] Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. Videoagent: Long-form video understanding with large language model as agent, 2024. 5
- [43] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. 6, 8
- [44] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa:next phase of question-answering to explaining temporal actions, 2021. 8
- [45] Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. In *European Conference on Computer Vision*, pages 131–147. Springer, 2025. 1
- [46] Jingkan Yang, Yuhao Dong, Shuai Liu, Bo Li, Ziyue Wang, Chencheng Jiang, Haoran Tan, Jiamu Kang, Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. Octopus: Embodied vision-language programmer from environmental feedback, 2023. 8
- [47] Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos. *arXiv preprint arXiv:2304.11968*, 2023. 3
- [48] 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. 6, 8
- [49] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6281–6290, 2019. 3, 4
- [50] Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, and Gedas Bertasius. A simple llm framework for long-range video question-answering, 2024. 1, 4, 5
- [51] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6586–6597, 2023. 5
- [52] Gengze Zhou, Yicong Hong, and Qi Wu. Navgpt: Explicit reasoning in vision-and-language navigation with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7641–7649, 2024. 5