

C^2 ROPE: Causal Continuous Rotary Positional Encoding for 3D Large Multimodal-Models Reasoning

Guanting Ye¹, Qiyan Zhao⁵, Wenhao Yu², Xiaofeng Zhang⁵, Jianmin Ji³, Yanyong Zhang⁴, and Ka-Veng Yuen^{1,*}

Abstract—Recent advances in 3D Large Multimodal Models (LMMs) built on Large Language Models (LLMs) have established the alignment of 3D visual features with LLM representations as the dominant paradigm. However, the inherited Rotary Positional Embedding (RoPE) introduces limitations for multimodal processing. Specifically, applying 1D temporal positional indices disrupts the continuity of visual features along the column dimension, resulting in spatial locality loss. Moreover, RoPE follows the prior that temporally closer image tokens are more causally related, leading to long-term decay in attention allocation and causing the model to progressively neglect earlier visual tokens as the sequence length increases. To address these issues, we propose C^2 RoPE, an improved RoPE that explicitly models local spatial Continuity and spatial Causal relationships for visual processing. C^2 RoPE introduces a spatio-temporal continuous positional embedding mechanism for visual tokens. It first integrates 1D temporal positions with Cartesian-based spatial coordinates to construct a triplet hybrid positional index, and then employs a frequency allocation strategy to encode spatio-temporal positional information across the three index components. Additionally, we introduce Chebyshev Causal Masking, which determines causal dependencies by computing the Chebyshev distance of image tokens in 2D space. Evaluation results across various benchmarks, including 3D scene reasoning and 3D visual question answering, demonstrate C^2 RoPE’s effectiveness. The code is available at <https://github.com/ErikZ719/C2RoPE>.

I. INTRODUCTION

Building 3D Large Multimodal Models (LMMs) [1], [2] has emerged as a promising research direction with significant potential for applications such as autonomous robotics and navigation systems. With the rapid advancement of Large Language Models (LLMs) [3], researchers are increasingly focusing on leveraging their impressive generation capabilities to enhance the understanding of 3D scene [4], [5].

Most mainstream 3D LMMs adopt a visual information injection paradigm similar to that of 2D LMMs [6]–[9], aligning visual representations with LLM embedding spaces

This work is funded by the Science and Technology Development Fund, Macau SAR under Research Grant 001/2024/SKL, the Research Committee of University of Macau under Research Grant CPG2026-00033-FST and the Guangdong-Hong Kong-Macau Joint Laboratory Program under Grant 2020B1212030009.

¹State Key Laboratory of Internet of Things for Smart City, University of Macau (UM), Macau, China

²Institute of Advanced Technology, University of Science and Technology of China (USTC), Hefei, China

³School of Computer Science and Technology, USTC, Hefei, China

⁴School of Artificial Intelligence and Data Science, USTC, Hefei, China

⁵Department of Automation, Shanghai Jiaotong University (SJTU), Shanghai, China

* Corresponding author. kv-yuen@um.edu.mo

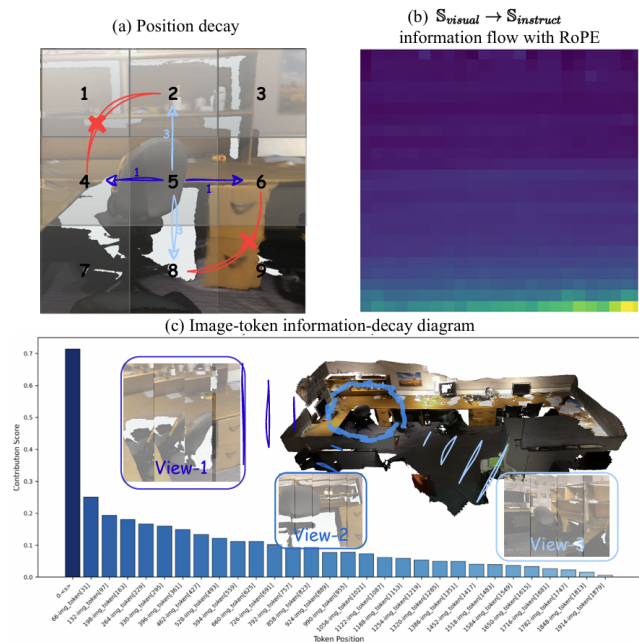


Fig. 1. (a) illustrates RoPE’s raster-scan positional index assignment; (b) visualizes the information flow from image tokens to instruction tokens, where flows from each image token are aggregated and arranged according to their positions in 2D space. The statistics are averaged over a 3K VQA subset. (c) shows the quantitative analysis of information flow from visual tokens to output tokens.

to enable cross-modal reasoning and instruction-conditioned policy synthesis for robotic navigation, manipulation, and human-robot interaction [10]–[12]. To construct effective 3D visual representations, some approaches project 3D points into 2D visual representations [13], [14], while others directly encode 3D features [15], [16]. More recently, LLaVA-3D [1] has innovatively embedded point cloud features into multi-view 2D image tokens, resulting in 3D perception-enhanced visual tokens. In these models, visual features are typically encoded as tokens and concatenated with text tokens to create a multi-modal input sequence, which is then processed autoregressively by LLMs. While this paradigm allows the 3D LMMs to be more compatible with LLM architectures, the implications of inheriting intrinsic mechanisms originally designed for natural language processing remain underexplored.

Recent research on information flow provides an intuitive interpretability approach for analyzing the intrinsic mechanisms of black-box models [17], [18]. Building on this, prior

studies [19]–[21] have shown through image-to-instruction information flow analysis that Rotary Position Encoding (RoPE) [22], a core component for positional modeling in LLMs, limits effective interaction between images and instructions in 2D LLMs. This raises the question:

For 3D LLMs, does inheriting RoPE from LLMs also introduce limitations?

In RoPE, image tokens are assigned unique temporal positional indices following a raster-scan order (left to right, top to bottom). However, this indexing scheme disrupts the continuity of visual features along the vertical (column) dimension. As illustrated in Fig. 1 (a), the positional indices of spatially adjacent image tokens are not continuous, a phenomenon we refer to as **spatial locality loss**.

Moreover, since RoPE treats temporally closer tokens as more causally relevant, it induces a long-term decay in attention allocation. To investigate information interaction under long-term decay, we visualize the image-to-instruction information flow. As shown in Fig. 1 (b), although visual tokens in 3D LLMs carry rich 3D information, only a small subset near the end of the sequence receives high attention, while most contribute only sparse information flow. We refer to this phenomenon as **visual tokens neglect**. Furthermore, Fig. 1 (c) shows the quantitative analysis of information flow from image tokens to output tokens, revealing that this neglect persists during response generation and becomes more pronounced as multi-view images further extend the sequence. A detailed discussion of spatial locality loss and visual tokens neglect is provided in Section 3.

Based on the above analysis, we propose a novel positional encoding method, **C^2 RoPE**, designed to preserve the local spatial Continuity of visual tokens and capture spatial Causal relationships. To achieve this, we introduce a spatio-temporal continuous positional modeling mechanism. First, the 2D image is projected onto a Cartesian coordinate system to obtain each token’s spatial coordinates (x, y) . These coordinates are then combined with RoPE’s original 1D temporal index m to form a triplet hybrid positional index (m, x, y) . Finally, we encode spatio-temporal information by assigning distinct frequency bands to each component in (m, x, y) . Additionally, we introduce Chebyshev Causal Masking, which determines the causal relationships of visual tokens based on their relative positional distance from the origin in the Cartesian coordinate system. This masking is applied to the self-attention matrix during decoding to alleviate image token neglect caused by long-term decay.

To fully evaluate the performance of C^2 RoPE, we conduct experiments on benchmarks including 3D scene reasoning and 3D visual question answering, demonstrating consistent performance improvements over the baselines. Compared to LLaVA-3D, our method achieves an improvement of +4.3 on EM@1 on ScanQA, and improvements of +1.2 on EM@1 and +1.2 on EM@R on SQA3D. Our main contributions are summarized as follows:

- This paper conducts an in-depth analysis of the limitations of applying RoPE in 3D LLMs, primarily manifested as spatial locality loss and visual tokens

neglect.

- Motivated by our analysis, we propose C^2 RoPE, which incorporates two key designs: a spatio-temporal continuous positional embedding mechanism and a Chebyshev causal masking strategy.
- Experiments on multiple benchmarks and across different baselines support the efficacy of our design.

II. RELATED WORK

A. 3D Large Multimodal Models

Large Language Models (LLMs) have made significant progress in recent years, demonstrating impressive performance across various scenarios [23]–[27]. These achievements have driven a growing number of researchers to develop Large Multimodal Models (LMMs) with 3D scene understanding capabilities [8], [28]–[31]. Current 3D LMMs typically use a pre-trained LLMs [32] as the decoder, and then achieve cross-modal feature alignment by projecting image representations into the LLM space and jointly decoding them with text tokens. Built on this basic architecture and the cross-modal alignment paradigm, subsequent research has primarily focused on constructing more effective 3D visual representations. For instance, ScanReason [33], PointLLM [16] and Li3da [4] leverage visual encoders to directly encode 3D point clouds features. Video-3D LLM [34], Scene-LLM [35], and 3D LLM [13] extract 3D visual representations from multi-view image inputs. The recently proposed LLaVA-3D [1] embeds 3D point cloud features into multi-view image tokens, resulting in 2D image tokens enhanced with 3D awareness. However, these 3D LMMs inherits the LLM’s positional modeling scheme, which can lead to issues such as “spatial locality loss” or the “image tokens neglect” problem.

B. Position Encoding Schemes in LMMs

Most 3D LMMs use position encoding methods inherited from LLMs, primarily RoPE [22], [36], [37]. In their approach, images are flattened into a 1D token sequence, with positional indices assigned to each image token in a raster-scan order. Subsequently, RoPE introduces complex rotations into the token embedding space. Compared to absolute [38] or relative position encoding [39], RoPE provides a smoother inductive bias for interactions between nearby and distant tokens. Subsequent works have proposed corresponding improvements to address the limitations of RoPE in specific domains. HoPE [40] proposes a hybrid frequency allocation strategy to improve the long-context capabilities. CoMemo [41] and MiniCPM-V [42] proposed different RoPE variants to alleviate remote decay in DHR scenarios. CCA [19] and MCA [20] proposed different heuristic positional index assignment strategies to mitigate hallucinations caused by long-term decay in 2D LLMs. To enhance multimodal processing, VideoRoPE [43] and M-RoPE [7] decomposing positional embedding into 1D, 2D, and 3D components and set different frequency distributions, effectively improving the performance of LMMs in video scenes. Although the aforementioned RoPE improvements

have enhanced performance in various scenarios, research has largely focused on 2D LMMs. Positional encoding in 3D LMMs remains underexplored. To the best of our knowledge, this work is the first to provide an in-depth analysis of RoPE’s limitations in the context of 3D LMMs.

III. ANALYSIS AND MOTIVATION

In this section, we further examine the spatial locality loss and visual tokens neglect, and conduct pilot experiments to investigate their impact. We begin by introducing the architecture of LLaVA-3D [1], the 3D LMMs we use as our baseline, and the mechanism of RoPE. Next, we elaborate on the spatial locality loss induced by RoPE’s positional indices, which fail to encode information along the column dimension. Finally, we present visualizations of image-to-instruction information flow, along with a quantitative analysis of information flow from image tokens to output token. These results reveal the image tokens neglect phenomenon. The above findings serve as the motivation and foundation for our proposed design.

A. Preliminary

3D LMMs: We adopt LLaVA-3D [1] as baseline, which consists of a pretrained image encoder F_v , a large language model F_t , and a projector f . Given multi-view 2D images $I_v \in \mathbb{R}^{V \times 3 \times H \times W}$, where V denotes the number of views. The image encoder F_v processes I_v into a multi-view 2D patch feature sequence $\mathbb{S}_{2d} \in \mathbb{R}^{V \times d \times h \times w}$, where d denotes the embedding dimension. LLaVA-3D then injects 3D positional embeddings $P_{3d} \in \mathbb{R}^{V \times 3 \times h \times w}$, encoded from point cloud features using a 2-layer MLP, into \mathbb{S}_{2d} , resulting in a 3D geometry-aware patch feature sequence $\mathbb{S}_{3d} = \mathbb{S}_{2d} + P_{3d}$. $\mathbb{S}_{3d} \in \mathbb{R}^{V \times d \times h \times w}$ are then sequentially sent into the projector f to obtain the final visual token sequence $\mathbb{S}_v = f(\mathbb{S}_{3d}) = \{w_m\}_{m=1}^v$. The instruction text I_t is encoded by F_t into the text tokens sequence $\mathbb{S}_t = F_t(I_t) = \{w_m\}_{m=1}^t$. Here, v and t denote the lengths of the visual and textual token sequences, respectively. Both \mathbb{S}_v and \mathbb{S}_t share the same embedding dimension d , and are concatenated to form the final multimodal input tokens sequence $\mathbb{S} = \text{Concat}(\mathbb{S}_v, \mathbb{S}_t) = \{w_m, m \in [1, v + t]\}$.

RoPE in LLaVA-3D: LLaVA-3D inherits RoPE from LLMs to encode the positional information of query and key tokens. Given a query token Q and a key token K with d dimensions, which are derived from \mathbb{S} . RoPE divides the dimension d into $d/2$ pairs, with each pair assigned predefined sinusoidal function values given by $\{\theta_i = 10000^{-2(i-1)/d}\}$, $i \in (1, 2, \dots, d/2)$. Subsequently, the rotation matrix $r^{(i)}$ is constructed as:

$$r^{(i)} = \begin{pmatrix} \cos(\theta_i) & -\sin(\theta_i) \\ \sin(\theta_i) & \cos(\theta_i) \end{pmatrix} \quad (1)$$

Each pairwise rotation matrix $r^{(i)}$ is concatenated along the diagonal to form the overall rotation matrix, denoted as $R_m = \text{diag}(r^{(1)}, \dots, r^{(d/2)})$, where m denotes the position index of each input token w_m . RoPE then applies rotation

matrix R_m to self-attention $A_{n,m}$ between the n -th query Q_n and m -th key K_m as:

$$A_{n,m} = \text{softmax} \left(\frac{(Q_n R_n)(K_m R_m)^T}{\sqrt{d}} \right) = \text{softmax} \left(\frac{Q_n R_{n-m} K_m^T}{\sqrt{d}} \right) \quad (2)$$

The rotation matrix R_{n-m} is formally given as follows:

$$R_{n-m} = \begin{pmatrix} \cos \theta_1(n-m) & -\sin \theta_1(n-m) & \dots & 0 & 0 \\ \sin \theta_1(n-m) & \cos \theta_1(n-m) & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & \cos \theta_{d/2}(n-m) & -\sin \theta_{d/2}(n-m) \\ 0 & 0 & \dots & \sin \theta_{d/2}(n-m) & \cos \theta_{d/2}(n-m) \end{pmatrix} \quad (3)$$

B. Spatial Locality Loss

In vanilla RoPE, the input 2D image I_v is flattened into the 1D tokens sequence $\mathbb{S}_v = \{w_m\}_{m=1}^v$, where m denotes the positional index of image tokens, which is assigned in a raster-scan order as follows:

$$\begin{bmatrix} 1 & 2 & \dots & \sqrt{v}-1 & \sqrt{v} \\ \sqrt{v}+1 & \sqrt{v}+2 & \dots & 2\sqrt{v}-1 & 2\sqrt{v} \\ & & \ddots & & \\ \vdots & \vdots & \frac{v}{2} - \frac{\sqrt{v}}{2} & \frac{v}{2} - \frac{\sqrt{v}}{2} + 1 & \vdots \\ \vdots & \vdots & \frac{v}{2} + \frac{\sqrt{v}}{2} & \frac{v}{2} + \frac{\sqrt{v}}{2} + 1 & \vdots \\ v-2\sqrt{v}+1 & v-2\sqrt{v}+2 & \dots & v-\sqrt{v}-1 & v-\sqrt{v} \\ v-\sqrt{v}+1 & v-\sqrt{v}+2 & \dots & v-1 & v \end{bmatrix} \quad (4)$$

Such image positional indices consider only the temporal positions of image tokens in the 1D input sequence while neglecting their positions in the 2D spatial domain. Specifically, although the sequential ordering of RoPE positional indices implicitly preserves the continuity of image tokens along the row dimension, the row-by-row alignment results in a spatial locality loss of visual features along the column dimension. In Fig. 1 (a), we illustrate this spatial locality loss more intuitively: along the column dimension, the positional indices of locally adjacent image tokens are not continuous.

C. Image Tokens Neglect

In RoPE, the rotation matrix R_{n-m} incorporates the relative positional distance information ($n-m$) between each query-key pair. This paradigm inherits the long-term decay property of sinusoidal encoding, where the decrease of $A_{n,m}$ as the $n-m$ increases. This prior is originally designed for natural language modeling, assuming temporally closer tokens are more causally related. In Fig. 1 (b), we present the visualization of image-to-instruction information flow to examine this property. We observe that attention is significantly concentrated on a small subset of visual tokens near the instruction tokens, which are spatially located in the lower-right region, while the majority are neglected.

We further analyze the quantitative results of information flow from image tokens to output tokens in Fig. 1 (c) and observe two important phenomena: (1) The initial tokens receive high attention. This is because, in causal self-attention, these tokens act as attention sinks to absorb redundant attention [17], [44]. (2) A large number of image tokens are neglected. The long-term decay of RoPE continues to affect autoregressive next-token prediction, causing visual tokens that are temporally distant from the output tokens

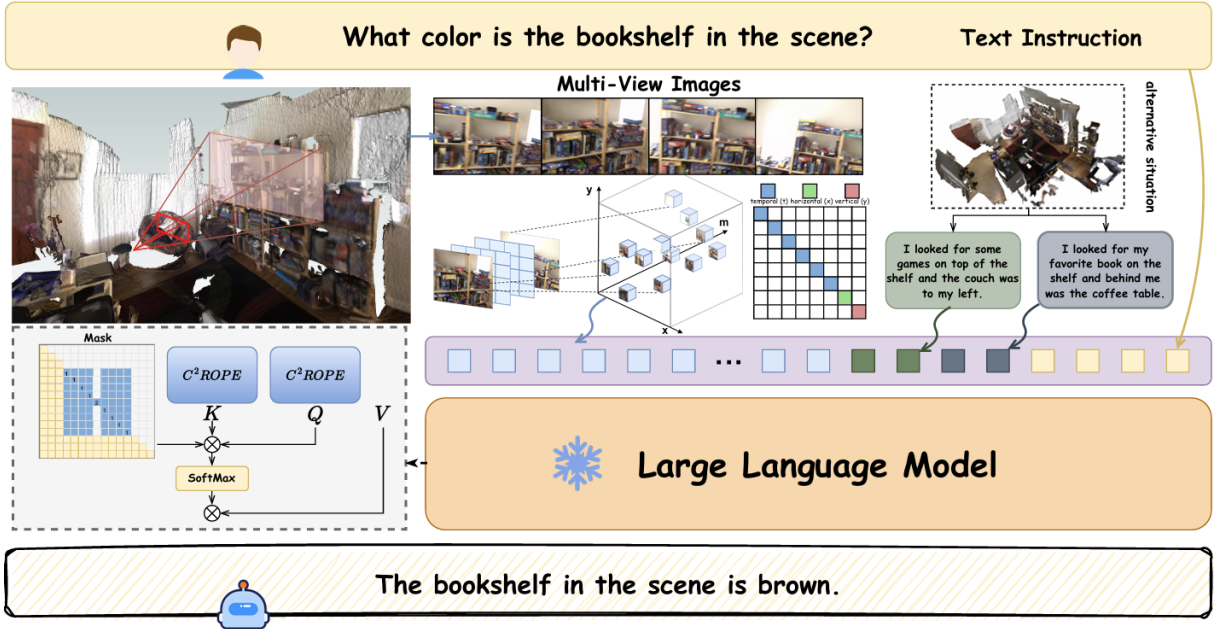


Fig. 2. Structure of C^2 RoPE. We integrate temporal indices and Cartesian spatial coordinates to form a triplet hybrid positional index that preserves local continuity of visual tokens. A frequency-based allocation encodes complementary spatio-temporal cues across the triplet, and Chebyshev Causal Masking enforces locality-aware causality by modulating attention with the 2D Chebyshev distance. This design mitigates long-range attention decay and improves 3D scene reasoning in LMMs.

to be considered weakly causal and thus assigned minimal attention. Moreover, LLaVA-3D introduces multi-view image inputs to represent 3D spatial information, introducing a critical trade-off: the growing sequence of image tokens exacerbates the image token neglect phenomenon.

IV. METHOD

To address these challenges, we propose C^2 RoPE, designed to enhance image scene information perception in LLaVA-3D. As shown in Fig. 4, C^2 RoPE consists of two key designs: (1) a spatio-temporal continuous positional embedding mechanism that explicitly models both the 1D temporal positions and 2D spatial positions of visual tokens to mitigate spatial locality loss; and (2) Chebyshev Causal Masking, which defines causal relationships based on the Chebyshev distance of image tokens in the Cartesian coordinate system, effectively alleviating RoPE's long-term decay and image token neglect. Details are presented in the following sections.

A. Spatio-temporal Continuous Positional Embedding Mechanism

Hybrid Image Positional Indexing: The analysis in Section 3.2 reveals that RoPE assigns positional indices to image tokens using a raster-scan order, which only represents their temporal positions in the input token sequence but fails to capture their positions in the spatial domain. Although the raster-scan positional indexing implicitly preserves the continuity of image tokens along the row dimension, it disrupts their continuity along the column dimension. As shown in Fig. 3 (d), we first project the image tokens onto a 2D Cartesian coordinate system to determine the spatial position of each image token. Specifically, the tokens at




Fig. 3. Long-distance position information attenuation of llava-3D, (a) represents the long-distance information attenuation in LLaVA, (b) (c) represent CCA-LLaVA [19] and MCA-LLaVA [20] to alleviate the decay of long-distance information by optimizing RoPE, (d) represents our method.

the image center are treated as the origin, with the axes aligned to rows and columns the positive directions extending rightward and upward, and the coordinates incrementing by 1. We then integrate each token's (x, y) coordinates with its original RoPE index m , forming a triplet hybrid image tokens positional index (m, x, y) . The final positional index of image tokens is as follows:

$$\begin{pmatrix} (1, 1 - \frac{\sqrt{e}}{e}, \frac{\sqrt{e}}{e} - 1) & (2, 2 - \frac{\sqrt{e}}{e}, \frac{\sqrt{e}}{e} - 1) & \dots & (\sqrt{e} - 1, \frac{\sqrt{e}}{e} - 2, \frac{\sqrt{e}}{e} - 1) & (\sqrt{e}, \frac{\sqrt{e}}{e} - 1, \frac{\sqrt{e}}{e} - 1) \\ (\sqrt{e} + 1, 1 - \frac{\sqrt{e}}{e}, \frac{\sqrt{e}}{e} - 2) & (\sqrt{e} + 2, 2 - \frac{\sqrt{e}}{e}, \frac{\sqrt{e}}{e} - 2) & \dots & (2\sqrt{e} - 1, \frac{\sqrt{e}}{e} - 2, \frac{\sqrt{e}}{e} - 2) & (2\sqrt{e}, \frac{\sqrt{e}}{e} - 1, \frac{\sqrt{e}}{e} - 2) \\ \vdots & \vdots & \dots & \vdots & \vdots \\ (\frac{e-2\sqrt{e}+1}{e}, \frac{e-2\sqrt{e}-2}{e}, \frac{e-2\sqrt{e}}{e}) & (\frac{e-2\sqrt{e}+2}{e}, \frac{e-2\sqrt{e}-2}{e}, \frac{e-2\sqrt{e}}{e}) & \dots & (\frac{e-\sqrt{e}-1}{e}, \frac{e-2\sqrt{e}-2}{e}, \frac{e-2\sqrt{e}}{e}) & (\frac{e-\sqrt{e}}{e}, \frac{e-2\sqrt{e}-1}{e}, \frac{e-2\sqrt{e}}{e}) \\ (\frac{e-\sqrt{e}+1}{e}, 1 - \frac{\sqrt{e}}{e}, 1 - \frac{\sqrt{e}}{e}) & (\frac{e-\sqrt{e}+2}{e}, 2 - \frac{\sqrt{e}}{e}, 2 - \frac{\sqrt{e}}{e}) & \dots & (\frac{e-1}{e}, \frac{e-2\sqrt{e}-1}{e}, \frac{e-2\sqrt{e}}{e}) & (\frac{e}{e}, \frac{e-2\sqrt{e}-1}{e}, \frac{e-2\sqrt{e}}{e}) \end{pmatrix} \quad (5)$$

To compare our design, we illustrate the positional index assignment of RoPE and its variants in Fig. 3 (a)-(b). Compared to these methods, our approach explicitly models both the temporal and spatial positional information of image tokens. Notably, the temporal component of our

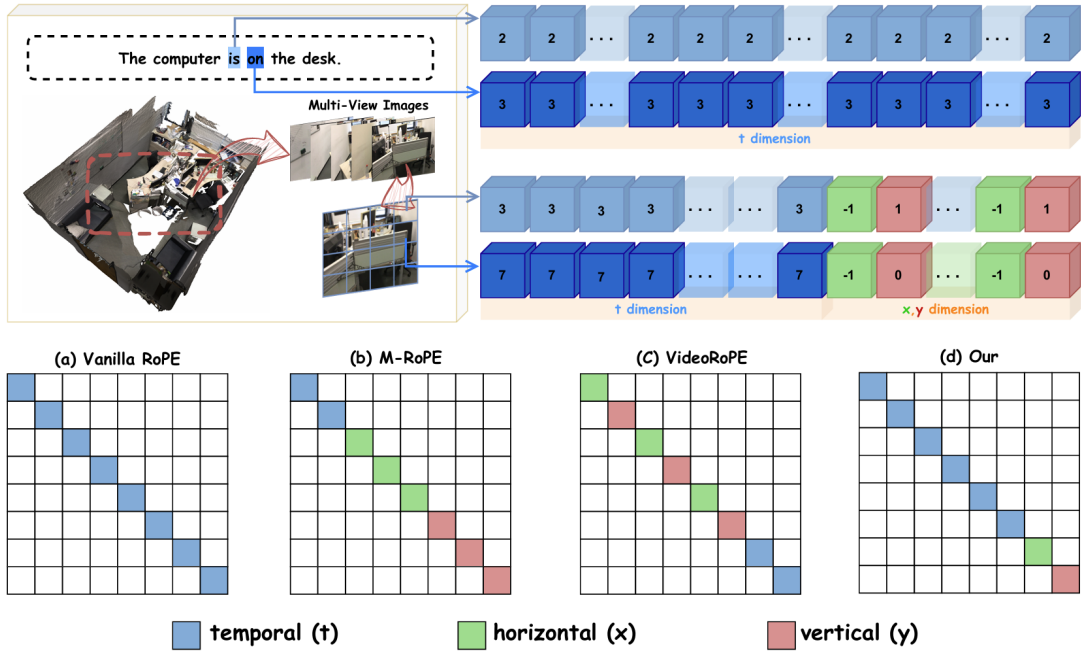


Fig. 4. The position embeddings of adjacent tokens for Vanilla RoPE (left), frame-aligned visual tokens for M-RoPE (center-left), interleaved spatio-temporal tokens for VideoRoPE (center-right), and our C^2 RoPE (right) with triplet hybrid indexing and Chebyshev causal masking.

TABLE I

PERFORMANCE COMPARISON ON SCANQA AND SQA3D BENCHMARKS. "EM@1" REPRESENTS TOP-1 EXACT MATCH, "B-4" REPRESENTS "BLEU-4", "MET" REPRESENTS "METEOR", "RGE" REPRESENTS "ROUGE", "EM@R" REPRESENTS REFINED EM1.

Methods	ScanQA					SQA3D (test)			ScanRefer	
	EM@1↑	B-4↑	MET↑	RGE↑	CIDEr↑	EM@1↑	EM@R↑	Acc@0.25↑	Acc@0.5↑	
<i>Expert Models</i>										
ScanQA [30]	21.1	10.1	13.1	33.3	64.9	-	-	-	-	
3D-VLP [45]	21.1	11.2	13.5	34.5	-	-	-	-	-	
3D-VisTA [46]	22.4	-	13.9	35.7	-	-	-	-	-	
<i>2D LLMs</i>										
InternV12-8B [47]	-	3.3	14.5	34.3	62.5	33.0	45.3	-	-	
Qwen2-VL-7B [7]	-	3.0	11.4	29.3	53.9	40.7	46.7	-	-	
LLaVA-Video-7B [48]	-	3.1	17.7	44.6	88.7	48.5	-	-	-	
<i>3D LLMs</i>										
Chat-3D [15]	-	6.4	11.9	28.5	53.2	-	-	35.9	30.4	
3D-LLM [13]	20.5	12.0	14.5	35.7	69.4	-	-	30.3	-	
Scene-LLM [35]	27.2	12.0	16.6	40.0	80.0	54.2	-	-	-	
LL3DA [4]	-	-	15.9	35.9	80.0	53.6	-	-	-	
LEO [31]	24.5	13.2	20.0	49.2	101.4	50.0	-	-	-	
ChatScene [49]	21.6	14.3	18.0	41.6	87.7	54.6	57.5	55.5	50.2	
Grounded 3D-LLM [50]	-	13.4	-	72.7	-	-	-	47.9	44.1	
Video-3D-LLM [34]	30.1	16.4	20.0	49.3	102.1	58.6	-	-	-	
Ross3D [5]	30.8	17.9	20.9	50.7	107.0	63.0	65.7	-	-	
AdaToken-3D [51]	30.0	13.3	17.6	44.7	91.8	54.8	-	-	-	
LLaVA-3D [1]	27.0	14.5	20.7	50.1	91.7	55.6	53.1	50.1	42.7	
Ours	31.3	23.0	34.1	52.6	109.8	56.8	54.3	50.5	42.9	

positional index reflects the ordering of image tokens within the input sequence, rather than the temporal correlation between adjacent video frames [40], [43].

Frequency Allocation Strategy: After constructing the triplet hybrid image positional index (m, x, y) , two further considerations arise:

Question1: How to ensure compatibility with the one-dimensional positional indices of text tokens?

In our positional index (m, x, y) , we retain RoPE's original index m for image tokens to maintain the absolute positions of image and text tokens in the original input

sequence. Additionally, we still use the temporal position m of image and text tokens for relative position calculation, ensuring consistency with the original RoPE when modeling relative positional dependencies. These designs remove the need for any additional processing of text token positional indices.

Question2: How to encode temporal and spatial information in hybrid image positional index?

Advanced video positional encoding methods, such as M-RoPE [7] and VideoRoPE [43], encode multi-dimensional information by assigning distinct frequency bands to different

components of the positional index. As shown in Fig. 4, RoPE allocates all 128 frequency dimensions to the temporal index. In contrast, M-RoPE assigns lower-frequency bands to the temporal component and higher-frequency bands to the horizontal and vertical components. VideoRoPE adopts the opposite strategy, prioritizing higher-frequency bands for the temporal component while alternately allocating lower-frequency bands to the horizontal and vertical components.

Notably, the temporal component in previous methods is defined at the image level, whereas our approach operates at the token level. For an image token with positional index (m, x, y) and a rotation matrix of dimension $d = 128$, we interleave x and y in the last 32 dimensions while assigning the remaining 96 dimensions to m . This design is driven by two considerations: (1) Lower dimensions correspond to higher frequencies and are highly sensitive to variations. Assigning high frequencies to positional components x, y may cause the model to overfocus on spatial positional changes of image tokens, thereby disrupting the LLMs’ well-trained temporal positional dependencies [52]; (2) RoPE induces a massive values pattern in query and key tokens, which is essential for capturing the contextual semantics of input tokens [53]. To preserve this property, we maintain a broad range of frequency bands aligned with RoPE in our allocation strategy.

B. Chebyshev Causal Masking

Standard RoPE implementations in language modeling assume that temporally closer tokens in text sequences are more causally related. The information flow analysis presented in Section 3.4 reveals that this assumption leads to long-term decay in attention allocation and the neglect of image tokens. Recent studies [19], [20] indicate that visual information is inherently structured across a two-dimensional spatial domain, where neighboring image tokens typically exhibit stronger causal relationships. Motivated by this observation, we propose the Chebyshev Causal Masking to explicitly encode this spatial causal relationships. Leveraging the spatial symmetry of images, we set the origin of the Cartesian coordinate system positioned at the image center as the reference point for modeling long-term decay. The causal correlation between tokens is determined by their Chebyshev distance from this origin: the farther a token is, the stronger its attention decay. Tokens sharing the same Chebyshev distance are grouped as correlated within visual self-attention. Fig. 3 illustrates the key differences between our approach and prior methods: we directly modify the default causal attention masking based on the prior of causal relationships in images. Additionally, we adopt the triplet image positional indexing described in Section 4.1, which preserves the continuity of image tokens in both temporal and spatial dimensions.

V. EXPERIMENTS

Experiment models This paper adopts LLaVA-3D-7B [1] as the baseline model and follows its two-stage training process, with each subsequent stage building upon the weights



Fig. 5. The case study is conducted on samples from SQA3D, where we provide identical model inputs to examine the generated responses.

learned in the previous stage. The number of views per input image is 16, and each image is encoded by a pre-trained image encoder into 576 2D image patches [54], resulting in a total of 9,216 patches. Similar to mainstream open-source LMMs, LLaVA-3D is built upon the powerful generative capabilities of LLMs. It achieves multimodal understanding by aligning visual features extracted by the CLIP encoder with the text embeddings of the LMMs. The key design of LLaVA-3D lies in endowing the model with 3D spatial understanding capabilities by injecting 3D information into the 2D visual tokens.

Implementation Details All experiments are executed on a setup comprising $4 \times$ H20 GPUs. Our 3D evaluation benchmark primarily relies on the Scannet Datasets [30], where we have selected 32 scenarios and rigorously tested thousands of QA pairs within each benchmark. In our ablation studies, all experimental settings and test sets are kept consistent to ensure fair comparisons. We plan to disclose the specific scenarios selected later to ensure the transparency and objectivity of our experimental findings. For multi-view inputs, we apply the same masking strategy and frequency-allocation policy uniformly to every view.

Benchmarks To comprehensively evaluate the 3D scene understanding capability of the proposed method, we conduct comprehensive experiments on two established 3D scene reasoning benchmarks: The SQA3D [55] benchmark enriches the evaluation by incorporating dynamic 3D scenes through 19K GPT-4-generated questions. The ScanQA [56] dataset provides 33.4K high-quality, human-annotated question to evaluate spatial relationships and navigation planning skills.

A. Comparison with State-of-the-Art Methods

As reported in Tab. I, compared to the baseline LLaVA-3D, our method achieves consistent improvements across all five evaluation metrics on ScanQA, with gains of +4.3 on EM@1, +8.5 on B-4, +13.4 on MET, +2.5 on RGE, and +18.1 on CIDEr. On the SQA3D test set, our method achieves improvements of +1.2 on EM@1 and +1.2 on EM@R, respectively.

Benefiting from the proposed spatio-temporal continuous positional embedding mechanism and Chebyshev causal

TABLE II

ABLATION STUDY ON THE SQA3D SUBSET, WHERE WE COMPARE C²RoPE WITH MULTIMODAL IMAGE POSITIONAL ENCODING METHODS CCA [19] AND MCA [20]. THE BASELINE LLaVA-3D ADOPTS THE DEFAULT RoPE.

Methods	SQA3D (test)		SQA3D (val)	
	EM@I ↑	EM@R ↑	EM@I ↑	EM@R ↑
LLaVA-3D	55.6	53.1	51.5*	54.5*
+MCA	55.5	53.7	50.9	54.7
+CCA	56.2	53.9	50.9	54.4
+C ² RoPE	56.8	54.3	51.6	55.7

masking, our method significantly enhances the model’s spatial reasoning and image perception capabilities. Specifically, our model achieves superior performance compared to expert models such as 3D-VLP [45], which are designed for specific tasks. This is because our model can leverage the powerful multimodal understanding capabilities of LMMs to significantly enhance performance on 3D tasks. Additionally, we present the performance of 2D LLMs on ScanQA and SQA3D. Compared to these models, our method achieves overwhelming performance advantages—for instance, it surpasses Qwen2-VL-7B by +20.0 on BLEU-4 and +22.7 on METEOR. Compared to current mainstream 3D LMMs, although our baseline does not achieve the best performance on ScanQA, it attains the state-of-the-art performance after incorporating the C²RoPE proposed in this paper, further demonstrating the effectiveness of our improvements.

We observe that LLaVA-3D underperforms compared to certain 3D LMMs, such as ChatScene and Ross3D, on the SQA3D benchmark. We analyze that this is because LLaVA-3D embeds point cloud features into 2D image tokens; while this enhances the 3D representation of image tokens, the model lacks richer 3D information as input for complex spatial reasoning in intricate scenes [55] (as illustrated in Figure 2). Although our method enhances the image perception of LLaVA-3D, it does not alter the way 3D information is modeled; therefore, its performance remains suboptimal compared to a few specialized 3D LMMs.

B. Ablation study

To mitigate long-term decay, CCA and MCA heuristically redirect instruction tokens to focus more on the central regions of images by reassigning different manually designed 1D positional indices to image tokens. In Fig. 3, we illustrate the long-term decay under RoPE, CCA, and MCA, respectively. In Fig. 3, we compare the information flow patterns of the baseline model and various improved image positional encoding methods. Our approach achieves a more balanced attention distribution across image tokens, enhancing the overall perception of visual information. Tab. II presents the quantitative results of the ablation study, where our method consistently achieves the best performance on both the validation and test subsets of SQA3D.

C. Case study

The Fig. 5 presents a case study in which both the baseline and our approach are provided with the same instruction

and multi-view scene images. In the displayed example, the baseline generates a response containing hallucinated content. In contrast, our method accurately perceives the visual information and produces a correct response.

VI. CONCLUSION

This paper reveals two key limitations of RoPE, inherited from LLMs, in the context of 3D LMMs: spatial locality loss and image token neglect. Through in-depth analysis, we explain the underlying causes: 1) RoPE assigns one-dimensional temporal positional indices to image tokens, disrupting their spatial continuity along the column dimension; 2) RoPE allocates attention based on the principle that “temporally closer image tokens are more causally related,” leading to long-term decay across image tokens. To address these issues, we propose C²RoPE, which incorporates two key designs: a spatio-temporal continuous positional embedding mechanism and Chebyshev causal masking. Comprehensive experiments validate the effectiveness of C²RoPE.

REFERENCES

- [1] C. Zhu, T. Wang, W. Zhang, J. Pang, and X. Liu, “Llava-3d: A simple yet effective pathway to empowering llms with 3d-awareness,” *arXiv preprint arXiv:2409.18125*, 2024.
- [2] J. Deng, T. He, L. Jiang, T. Wang, F. Dayoub, and I. Reid, “3d-llava: Towards generalist 3d llms with omni superpoint transformer,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 3772–3782.
- [3] A. Yang, B. Yang, B. Hui, B. Zheng, B. Yu, C. Zhou, C. Li, C. Li, D. Liu, *et al.*, “Qwen3 technical report,” *arXiv preprint arXiv:2505.09388*, 2025.
- [4] S. Chen, X. Chen, C. Zhang, M. Li, G. Yu, H. Fei, H. Zhu, J. Fan, and T. Chen, “L13da: Visual interactive instruction tuning for omni-3d understanding reasoning and planning,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 26 428–26 438.
- [5] H. Wang, Y. Zhao, T. Wang, H. Fan, X. Zhang, and Z. Zhang, “Ross3d: Reconstructive visual instruction tuning with 3d-awareness,” *arXiv preprint arXiv:2504.01901*, 2025.
- [6] H. Liu, C. Li, Y. Li, and Y. J. Lee, “Improved baselines with visual instruction tuning,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 26 296–26 306.
- [7] P. Wang, S. Bai, S. Tan, S. Wang, Z. Fan, J. Bai, K. Chen, X. Liu, J. Wang, W. Ge, *et al.*, “Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution,” *arXiv preprint arXiv:2409.12191*, 2024.
- [8] T. Zhang, X. Li, H. Fei, H. Yuan, S. Wu, S. Ji, C. C. Loy, and S. Yan, “Omg-llava: Bridging image-level, object-level, pixel-level reasoning and understanding,” *Advances in neural information processing systems*, vol. 37, pp. 71 737–71 767, 2024.
- [9] G. Team, R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican, *et al.*, “Gemini: a family of highly capable multimodal models,” *arXiv preprint arXiv:2312.11805*, 2023.
- [10] W. Yu, J. Peng, Q. Qiu, H. Wang, L. Zhang, and J. Ji, “Pathrl: An end-to-end path generation method for collision avoidance via deep reinforcement learning,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 9278–9284.
- [11] Y. Duan, H. Li, Y. Wu, W. Yu, X. Zhang, Y. Shen, J. Ji, and Y. Zhang, “Stdarm: Transferring visuomotor policies from static data training to dynamic robot manipulation,” *arXiv preprint arXiv:2504.18792*, 2025.
- [12] W. Yu, J. Peng, H. Yang, J. Zhang, Y. Duan, J. Ji, and Y. Zhang, “Ldp: A local diffusion planner for efficient robot navigation and collision avoidance,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 5466–5472.
- [13] Y. Hong, H. Zhen, P. Chen, S. Zheng, Y. Du, Z. Chen, and C. Gan, “3d-llm: Injecting the 3d world into large language models,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 20 482–20 494, 2023.

- [14] H. Zhi, P. Chen, J. Li, S. Ma, X. Sun, T. Xiang, Y. Lei, M. Tan, and C. Gan, "Lscenellm: Enhancing large 3d scene understanding using adaptive visual preferences," *arXiv preprint arXiv:2412.01292*, 2024.
- [15] Z. Wang, H. Huang, Y. Zhao, Z. Zhang, and Z. Zhao, "Chat-3d: Data-efficiently tuning large language model for universal dialogue of 3d scenes," *arXiv preprint arXiv:2308.08769*, 2023.
- [16] R. Xu, X. Wang, T. Wang, Y. Chen, J. Pang, and D. Lin, "Pointllm: Empowering large language models to understand point clouds," in *European Conference on Computer Vision*. Springer, 2024, pp. 131–147.
- [17] L. Wang, L. Li, D. Dai, D. Chen, H. Zhou, F. Meng, J. Zhou, and X. Sun, "Label words are anchors: An information flow perspective for understanding in-context learning," *ArXiv*, vol. abs/2305.14160, 2023.
- [18] L. Chen, H. Zhao, T. Liu, S. Bai, J. Lin, C. Zhou, and B. Chang, "An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models," in *European Conference on Computer Vision*. Springer, 2024, pp. 19–35.
- [19] Y. Xing, Y. Li, I. Laptev, and S. Lu, "Mitigating object hallucination via concentric causal attention," *Advances in neural information processing systems*, vol. 37, pp. 92012–92035, 2024.
- [20] Q. Zhao, X. Zhang, Y. Li, Y. Xing, X. Yuan, F. Tang, S. Fan, X. Chen, X. Zhang, and D. Wang, "Mca-llava: Manhattan causal attention for reducing hallucination in large vision-language models," *arXiv preprint arXiv:2507.09184*, 2025.
- [21] Y. Zhu, L. Tao, M. Dong, and C. Xu, "Mitigating object hallucinations in large vision-language models via attention calibration," *ArXiv*, vol. abs/2502.01969, 2025.
- [22] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu, "Roformer: Enhanced transformer with rotary position embedding," *Neurocomputing*, vol. 568, p. 127063, 2024.
- [23] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al., "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [24] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al., "Palm: Scaling language modeling with pathways," *Journal of Machine Learning Research*, vol. 24, no. 240, pp. 1–113, 2023.
- [25] W. Dai, J. Li, D. Li, A. Tiong, J. Zhao, W. Wang, B. Li, P. N. Fung, and S. Hoi, "Instructblip: Towards general-purpose vision-language models with instruction tuning," *Advances in neural information processing systems*, vol. 36, pp. 49250–49267, 2023.
- [26] Y. Li, Y. Du, K. Zhou, J. Wang, W. X. Zhao, and J.-R. Wen, "Evaluating object hallucination in large vision-language models," *arXiv preprint arXiv:2305.10355*, 2023.
- [27] X. Wang, S. Tan, M. Jin, W. Y. Wang, R. Panda, and Y. Shen, "Do larger language models imply better generalization? a pretraining scaling law for implicit reasoning," 2025.
- [28] J. Chen, D. Zhu, X. Shen, X. Li, Z. Liu, P. Zhang, R. Krishnamoorthi, V. Chandra, Y. Xiong, and M. Elhoseiny, "Minigpt-v2: large language model as a unified interface for vision-language multi-task learning," *arXiv preprint arXiv:2310.09478*, 2023.
- [29] Z. Chen, A. Gholami, M. Nießner, and A. X. Chang, "Scan2cap: Context-aware dense captioning in rgb-d scans," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 3193–3203.
- [30] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "Scannet: Richly-annotated 3d reconstructions of indoor scenes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5828–5839.
- [31] J. Huang, S. Yong, X. Ma, X. Linghu, P. Li, Y. Wang, Q. Li, S.-C. Zhu, B. Jia, and S. Huang, "An embodied generalist agent in 3d world," *arXiv preprint arXiv:2311.12871*, 2023.
- [32] A. M. Kassem, O. Mahmoud, N. Mireshghallah, H. Kim, Y. Tsvetkov, Y. Choi, S. Saad, and S. Rana, "Alpaca against vicuna: Using llms to uncover memorization of llms," *arXiv preprint arXiv:2403.04801*, 2024.
- [33] C. Zhu, T. Wang, W. Zhang, K. Chen, and X. Liu, "Scanreason: Empowering 3d visual grounding with reasoning capabilities," 2024.
- [34] D. Zheng, S. Huang, and L. Wang, "Video-3d llm: Learning position-aware video representation for 3d scene understanding," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 8995–9006.
- [35] R. Fu, J. Liu, X. Chen, Y. Nie, and W. Xiong, "Scene-llm: Extending language model for 3d visual understanding and reasoning," *arXiv preprint arXiv:2403.11401*, 2024.
- [36] M. Gao, J. Liu, M. Li, J. Xie, Q. Liu, B. Zhao, X. Chen, and H. Xiong, "Tc-llava: Rethinking the transfer from image to video understanding with temporal considerations," *arXiv preprint arXiv:2409.03206*, 2024.
- [37] P. Agrawal, S. Antoniak, E. B. Hanna, B. Bout, D. Chaplot, J. Chudnovsky, D. Costa, B. De Monicault, S. Garg, T. Gervet, et al., "Pixtral 12b," *arXiv preprint arXiv:2410.07073*, 2024.
- [38] A. Vaswani, N. M. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Neural Information Processing Systems*, 2017.
- [39] G. Ke, D. He, and T.-Y. Liu, "Rethinking positional encoding in language pre-training," in *International Conference on Learning Representations*, 2021.
- [40] H. Li, Y. Qin, B. Ou, L. Xu, and R. Xu, "Hope: Hybrid of position embedding for length generalization in vision-language models," *arXiv preprint arXiv:2505.20444*, 2025.
- [41] S. Liu, W. Su, X. Zhu, W. Wang, and J. Dai, "Comemo: Lvlms need image context with image memory," *arXiv preprint arXiv:2506.06279*, 2025.
- [42] Y. Yao, T. Yu, A. Zhang, C. Wang, J. Cui, H. Zhu, T. Cai, H. Li, W. Zhao, Z. He, Q.-A. Chen, H. Zhou, Z. Zou, H. Zhang, S. Hu, Z. Zheng, J. Zhou, J. Cai, X. Han, G. Zeng, D. Li, Z. Liu, and M. Sun, "Minicpm-v: A gpt-4v level mllm on your phone," *ArXiv*, vol. abs/2408.01800, 2024.
- [43] X. Wei, X. Liu, Y. Zang, X. Dong, P. Zhang, Y. Cao, J. Tong, H. Duan, Q. Guo, J. Wang, et al., "Videorope: What makes for good video rotary position embedding?" *arXiv preprint arXiv:2502.05173*, 2025.
- [44] X. Zhang, Y. Quan, C. Gu, C. Shen, X. Yuan, S. Yan, H. Cheng, K. Wu, and J. Ye, "Seeing clearly by layer two: Enhancing attention heads to alleviate hallucination in lvlms," *ArXiv*, vol. abs/2411.09968, 2024.
- [45] Z. Jin, M. Hayat, Y. Yang, Y. Guo, and Y. Lei, "Context-aware alignment and mutual masking for 3d-language pre-training," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 10984–10994.
- [46] Z. Zhu, X. Ma, Y. Chen, Z. Deng, S. Huang, and Q. Li, "3d-vista: Pre-trained transformer for 3d vision and text alignment," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 2911–2921.
- [47] O. Team, "Internvl2: Better than the best—expanding performance boundaries of open-source multimodal models with the progressive scaling strategy," 2024.
- [48] Y. Zhang, J. Wu, W. Li, B. Li, Z. Ma, Z. Liu, and C. Li, "Video instruction tuning with synthetic data," *arXiv preprint arXiv:2410.02713*, 2024.
- [49] J. Zhang, C. Xu, and B. Li, "Chatscene: Knowledge-enabled safety-critical scenario generation for autonomous vehicles," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15459–15469.
- [50] Y. Chen, S. Yang, H. Huang, T. Wang, R. Xu, R. Lyu, D. Lin, and J. Pang, "Grounded 3d-llm with referent tokens," *arXiv preprint arXiv:2405.10370*, 2024.
- [51] K. Zhang, X. Chen, and X. Zhang, "Adatoken-3d: Dynamic spatial gating for efficient 3d large multimodal-models reasoning," *arXiv preprint arXiv:2505.12782*, 2025.
- [52] F. Barbero, A. Vitvitskiy, C. Perivolaropoulos, R. Pascanu, and P. Veličković, "Round and round we go! what makes rotary positional encodings useful?" in *The Thirteenth International Conference on Learning Representations*, 2025.
- [53] M. Jin, K. Mei, W. Xu, M. Sun, R. Tang, M. Du, Z. Liu, and Y. Zhang, "Massive values in self-attention modules are the key to contextual knowledge understanding," *arXiv preprint arXiv:2502.01563*, 2025.
- [54] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*, 2021.
- [55] X. Ma, S. Yong, Z. Zheng, Q. Li, Y. Liang, S.-C. Zhu, and S. Huang, "Sqa3d: Situated question answering in 3d scenes," *ArXiv*, vol. abs/2210.07474, 2022.
- [56] D. Azuma, T. Miyanishi, S. Kurita, and M. Kawanabe, "Scanqa: 3d question answering for spatial scene understanding," *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19107–19117, 2021.