

Towards Proprioception-Aware Embodied Planning for Dual-Arm Humanoid Robots

Boyu Li^{1,2,3}, Siyuan He⁴, Hang Xu⁴, Haoqi Yuan^{5,6},
 Yu Zang⁴, Liwei Hu⁴, Junpeng Yue⁵, Zhenxiong Jiang⁴, Pengbo Hu⁴,
 Börje F. Karlsson³, Dongbin Zhao^{1,2,3}, Yehui Tang^{4†}, Zongqing Lu^{5,6*}

¹SKL-MAIS, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³Beijing Academy of Artificial Intelligence, ⁴AgiBot

⁵School of Computer Science, Peking University, ⁶BeingBeyond

Abstract—In recent years, Multimodal Large Language Models (MLLMs) have demonstrated the ability to serve as high-level planners, enabling robots to follow complex human instructions. However, their effectiveness, especially in long-horizon tasks involving dual-arm humanoid robots, remains limited. This limitation arises from two main challenges: (i) the absence of simulation platforms that systematically support task evaluation and data collection for humanoid robots, and (ii) the insufficient embodiment awareness of current MLLMs, which hinders reasoning about dual-arm selection logic and body positions during planning. To address these issues, we present DualTHOR, a new dual-arm humanoid simulator, with continuous transition and a contingency mechanism. Building on this platform, we propose Proprio-MLLM, a model that enhances embodiment awareness by incorporating proprioceptive information with motion-based position embedding and a cross-spatial encoder. Experiments show that, while existing MLLMs struggle in this environment, Proprio-MLLM achieves an average improvement of 19.75% in planning performance. Our work provides both an essential simulation platform and an effective model to advance embodied intelligence in humanoid robotics. The code is available at <https://anonymous.4open.science/r/DualTHOR-5F3B/>.

I. INTRODUCTION

In recent years, leveraging Multimodal Large Language Models (MLLMs) as high-level planners to enable robots to execute complex human instructions has become a prominent research direction in Embodied AI [1]. To this end, the research community is actively pursuing a deep synergy between MLLMs and highly versatile dual-arm humanoid robots, aiming to unlock their vast potential in fine-grained [2], long-horizon [3], [4], and human-like

[†]Project Lead.

*Correspondence to Zongqing Lu <zongqing.lu@pku.edu.cn>.

This research was supported in part by NSFC under Grant 62450001, 62136008, Doubao Fund and the Suzhou Innovation and Entrepreneurship Leading Talents Programme - Innovation Leading Talent in Universities and Research Institutes with Grant No. ZXL2025310.

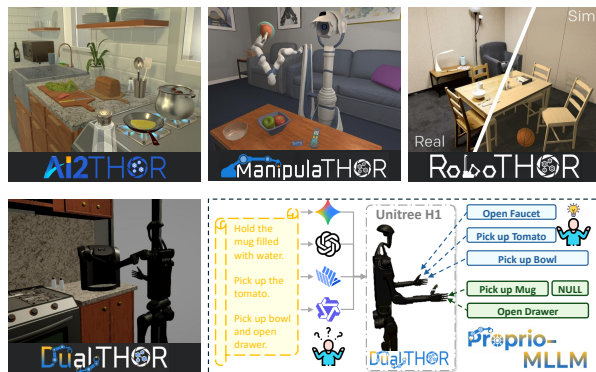


Fig. 1: **DualTHOR** is a novel simulator specifically tailored for dual-arm humanoid robots, while still preserving the diversity and realism of scenarios in previous AI2THOR series simulators. As current MLLMs have limited effectiveness in planning for dual-arm embodied tasks, we propose **Proprio-MLLM** to achieve proprioception-aware, embodiment-grounded planning.

collaborative tasks [5]. However, advancing this research frontier faces a fundamental bottleneck stemming from both data and models: first, a severe scarcity of high-level planning simulators for dual-arm humanoid robots that can support training data collection and evaluation of MLLMs; and second, the inherent capability limitations of existing MLLMs for dual-arm humanoid robot high-level planning tasks.

This high-level planning data gap arises from systemic limitations in the current benchmarks [6], [7], which fail to provide the long-horizon, visual-physically realistic humanoid interaction planning data required for MLLM training [8]. Specifically, this deficiency manifests in two aspects. First, existing benchmarks for long-horizon task planning [9], [10] are predominantly confined to wheeled robots or single-arm manipulators, as illustrated in the upper part of Fig. 1. Moreover, they often sacrifice

the physical realism of low-level control to focus on high-level logical evaluation. They rely on idealized, instantaneous state transitions [11] to ensure flawless skill sequencing, but this simplification produces data that is detached from physical reality, limiting the real-world applicability of MLLM planning policies and their sim-to-real transferability. Second, although some benchmarks support bimanual humanoid robots, they concentrate almost entirely on low-level control [12], [13], neglecting long-horizon planning challenges. Their evaluation scope is narrow, generating only fragmented data on isolated sub-skills such as picking and placing, without providing coherent task sequences necessary to assess performance in complex, multi-stage daily activities [14]. Consequently, there is an urgent need for a benchmark platform capable of producing high-quality humanoid datasets that integrate finer-grained low-level control with long-horizon planning, filling this critical gap in the field.

Beyond the data scarcity in the current benchmarks, the second fundamental challenge lies within the MLLMs themselves. Although these models show encouraging capabilities in long-horizon embodied planning tasks [15], their effectiveness is severely constrained by the lack of grounding in the dual-arm humanoid robots’ physical reality [16], as shown in the lower part of Fig. 1. By failing to incorporate robots’ embodiment, MLLMs often struggle with arm selection logic, body state adjustments (e.g., height changes), and the identification of coherent and physically plausible interaction points with objects [17]. For instance, sometimes the target object is located on the left side of the robot, but the left hand is already occupied based on the MLLM’s prior plan, while the right hand cannot reach it. Such conflicts typically necessitate re-planning for revised locomotion [18] and manipulation, significantly reducing task success rates within a reasonable number of steps. Also, MLLMs should consider the robots’ interaction range based on their joint configurations to find proper interaction points [11], as low-level skills must always operate within a constrained range to guarantee stable execution. These limitations emphasize the importance of enhancing MLLMs’ understanding of robot’s embodiment.

To solve the aforementioned challenges, we first introduce a dual-arm humanoid simulation platform, **DualTHOR**, to support data collection and task evaluation for MLLMs. We employ dual-arm humanoid robots as primary agents and design a set of dual-arm tasks in which the two arms can either execute distinct actions in parallel or collaborate to accomplish a single complex task. We optimize the control logic to ensure continuous robot’s states and environmental transitions, and introduce a contingency mechanism to further simulate real-world uncertainty. Building on this foundation, We further

propose **Proprio-MLLM**, an MLLM that incorporates proprioceptive information into dual-arm planning. We introduce a motion-based position embedding method and a cross-spatial encoder to enhance the model’s embodiment awareness and spatial reasoning abilities. Experimental results show that while existing MLLMs struggle with dual-arm planning tasks, Proprio-MLLM achieves an average improvement of 19.75% in planning performance, demonstrating more reliable logical and spatial reasoning capabilities.

In summary, our contributions encompass the following key advancements:

- 1) We propose a dual-arm humanoid robot simulator **DualTHOR** (based on AI2-THOR) and introduce a task suite for dual-arm planning. We create a new benchmark tailored for household dual-arm tasks, providing a standardized evaluation framework for future research.
- 2) By incorporating proprioceptive information, we propose a multimodal alignment large language model, **Proprio-MLLM**. We introduce a motion-based position embedding method and a cross-spatial encoder, increasing the model’s embodiment awareness and spatial reasoning in dual-arm tasks.
- 3) Our experimental results show that existing MLLMs have limited capabilities for dual-arm embodied planning tasks and Proprio-MLLM achieves an average improvement of 19.75% in planning performance.

II. RELATED WORK

A. Simulation Platforms for Interactive Learning

Simulation environments are indispensable for advancing interactive learning in embodied AI, offering controllable settings for agent training and evaluation. AI2-THOR [9] serves as a foundational framework, providing high-quality visual environments and interactive object dynamics that support a wide range of embodied AI tasks. Extensions such as RoboTHOR [6], ManipulaTHOR [10], and ProcTHOR [7] have enhanced the simulator by increasing task diversity, scene complexity, and object variety, yet the range of robot embodiments remains limited to single-arm and wheeled robots. HumanoidBench [12] and Isaac Gym [13] focus primarily on training low-level control policies for humanoid robots. Consequently, their scenes assets are limited for assessing long-horizon planning capabilities for MLLMs. To provide the necessary simulation capabilities along with a wide range of task categories and extended robot morphologies, we introduce our DualTHOR simulator.

TABLE I: A systematic comparison of DualTHOR and existing long-horizon household simulation platforms.

Simulator	Category	Agents	Transition	Action Control	Contingency
ThreeDWorld [19]	Household	Wheel	Discrete	Discrete	✗
iGibson [20]	Navigation	Wheel	Discrete	Discrete	✗
AI2-THOR [9]	Household	Wheel	Discrete	Discrete	✗
RoboThor [6]	Navigation	Wheel	Discrete	Discrete	✗
ManipulaThor [10]	Manipulation	Single Arm	Discrete	Discrete	✗
ProcThor [7]	Household	Wheel	Discrete	Discrete	✗
OmniGibson [21]	Multi-Domain	Single Arm & Wheel	Discrete	Discrete	✗
MoMa-Kitchen [11]	Manipulation	Single Arm	Discrete	Discrete	✗
DualTHOR (ours)	Household	Dual Arm	Continuous	Discrete & Continuous	✓

B. Multimodal Alignment for Large Language Models

Recent research on alignment techniques in MLLMs can be broadly categorized into deep fusion and shallow fusion methods. Deep fusion approaches, such as Flamingo [22] and NVLM [23], modify the LLM’s architecture by incorporating cross-attention layers and additional feedforward components to improve modality alignment. In contrast, shallow fusion methods, including LLaVA [24] and MotionGPT [25], map multimodal latent features into the LLM’s embedding space using either MLP-based mappings or cross-attention mechanisms. Shallow fusion methods offer greater computational efficiency, making them ideal for integrating proprioceptive information into MLLMs in a resource-efficient manner.

III. DUALTHOR

We introduce DualTHOR, a novel simulation platform for dual-arm humanoid robots extended from AI2-THOR. DualTHOR significantly advances embodied interaction complexity by introducing a new suite of long-horizon tasks designed specifically for bimanual humanoid robots. Crucially, and in contrast to existing platforms, it supports continuous, physically-grounded interactions and incorporates a stochastic contingency mechanism to enable MLLMs to develop and refine their re-planning abilities in response to execution errors, as shown in Table I. With these enhancements, DualTHOR provides a highly interactive and flexible platform for advancing the development of robust dual-arm embodied systems.

A. Overview of DualTHOR

Physics Engine. We build DualTHOR on the Unity engine to support physically realistic bimanual manipulation. Unity’s framework facilitates the parallel execution of coordinated arm actions and ensures smooth, continuous transitions through interpolation. This physical and visual fidelity supported by our platform is critical for perception-driven agents operating in the diverse and complex household scenarios. DualTHOR is a multi-camera suite designed to provide comprehensive situational awareness across diverse household scenes (see Fig. 2) while mitigating self-occlusion.

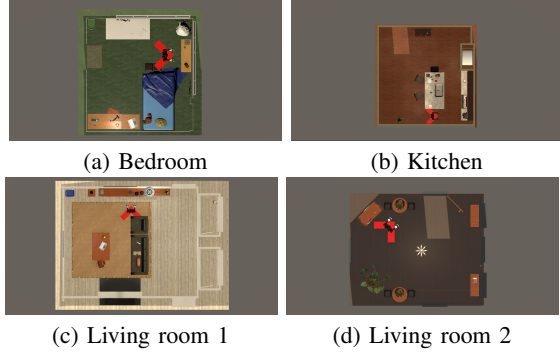


Fig. 2: **Example scenes of different rooms in DualTHOR.** The types and quantities of objects vary across rooms, and the humanoid robot is capable of interacting with all objects within each room.

Humanoid Robots. To support a broad range of interaction scenarios, DualTHOR integrates two distinct humanoid robots: the powerful and stable Unitree H1, and the dexterous and precise Agibot X1. This diversity enables research on tasks ranging from heavy-duty interactions (H1) to fine-motor control (X1). Crucially, the two robots feature different end-effectors—dexterous hands for H1 and parallel grippers for X1—allowing the study of different manipulation paradigms within a unified platform.

B. Task Categories

Tasks in DualTHOR are designed to rigorously evaluate bimanual capabilities. To this end, we introduce a novel categorization based on the required number of arms:

Dual-Arm Essential Tasks: Tasks that are physically impossible with a single arm. Examples include lifting heavy objects or coordinating to hold an object while opening an affordance-specific container.

Dual-Arm Optional Tasks: Tasks that can be performed with a single arm, but are enhanced through dual-arm execution, such as transporting two objects within constrained timesteps.

Single-Arm Tasks: Simple interaction tasks, largely adapted from AI2-THOR, that serve as basic plan tasks.

C. Low-Level Control

Low-level control in DualTHOR operates on a client-server model, where a Python API sends high-level commands to the Unity engine. For interaction actions, Unity queries a modular inverse kinematics (IK) service that uses OmniManip [26] to compute the required joint configurations. This solver is particularly suited for our platform due to its support for dexterous hand models, and the resulting trajectories are smoothed to ensure continuous, realistic motion.

A key aspect of our low-level control is the implementation of distinct IK models tailored to each robot’s specific architecture. The Agibot X1 adopts a decoupled approach, solving the IK for each arm independently using a simplified pose matrix (a 3×3 rotation matrix with a translation vector) relative to its base frame. In contrast, the Unitree H1 utilizes a whole-body coordination model that solves for both arms simultaneously. This model employs a complete 4×4 homogeneous transformation matrix to incorporate full posture information and integrates current joint states (angles and velocities) for dynamic optimization, with a specific focus on maintaining balance constraints during bimanual operations. This dual-model approach allows DualTHOR to faithfully simulate the distinct control paradigms of different real-world humanoid systems.

D. Contingency Mechanism

We introduce a stochastic contingency mechanism in DualTHOR. This system is designed to simulate real-world uncertainty, by mapping an action to a distribution of potential outcomes rather than a single guaranteed result. For example, when an agent attempts to pick up a “pourable” cup, there is an 80% probability of success, but also a 10% chance of the cup breaking and a 10% chance of its contents spilling (illustrated in Fig. 3). These stochastic outcomes, derived from object categories, allow the agent to move beyond rigid plans and require MLLMs to develop robust error-recovery and re-planning capabilities.

IV. METHOD

In this work, we propose **Proprio-MLLM**, a model for long-horizon dual-arm robot planning that tackles two key challenges: (i) integrating proprioceptive information into the planning process to resolve arm selection logic and body state changes, and (ii) enhancing spatial reasoning to ensure coherent and physically plausible interaction points with objects. To achieve this, we extend Qwen 2.5-VL-7B-Ins [27] with multi-modal motion inputs, enabling it to perceive and reason over the robot’s proprioceptive states. We introduce a motion-based position embedding and a cross-spatial encoder for proprioceptive information grounding. The overall architecture is shown in Fig. 4.




Fig. 3: **Example of picking up a “pourable” cup of coffee.** The possible results include success (80%), coffee spill (10%), and mug broken (10%). DualTHOR provides environmental feedback after the robot executes an action, enabling the evaluation of MLLM re-planning abilities.

A. Alignment Dataset Preparation

Motivated by MotionGPT [25], we use motion modal to help MLLMs understand the proprioceptive states of dual-arm humanoid robots. We build a robot motion-image-text dataset upon the HumanML3D dataset [28], which provides different human action samples for motion generation tasks. As the HumanML3D dataset provides humanoid motion in SMPL format [29], we re-target the human motion data into the H1, X1 robots and replay them in DualTHOR simulator to obtain corresponding first-person visual observations. After PHC-based filtering [16], we obtain 11k high-quality trajectories for training.

B. Motion Alignment Training

We separate the training scheme into three stages. First, we learn the motion tokenizer to align with the MLLM’s token-in-token-out format. Second, we fine-tune the MLLM to achieve motion, language and image feature alignment. Finally, we conduct instruction tuning to enable Proprio-MLLM to solve embodied dual-arm tasks.

Stage 1: Motion Tokenizer. We introduce a motion VQ-VAE [30] model, which consists of a motion encoder E , a motion decoder D , and a codebook $C = \{c_1, c_2, \dots, c_M\}$ of size M . We denote a robot motion sequence from alignment dataset as $m = \{m_1, m_2, \dots, m_T\} \in \mathbb{R}^{T \times d}$, where T is the sequence length and d is the feature dimension per frame. The encoder E projects m into a latent embedding $z(m)$, which is then discretized by quantization. Specifically, each embedding is mapped to its nearest entry in the learnable codebook C . The quantized vector e and its code index p are defined as:

$$e = c_p, \quad p = \arg \min_k \|z(m) - c_k\|_2. \quad (1)$$

The motion VQ-VAE is trained with three tailored loss: a reconstruction loss, a codebook loss and a commitment

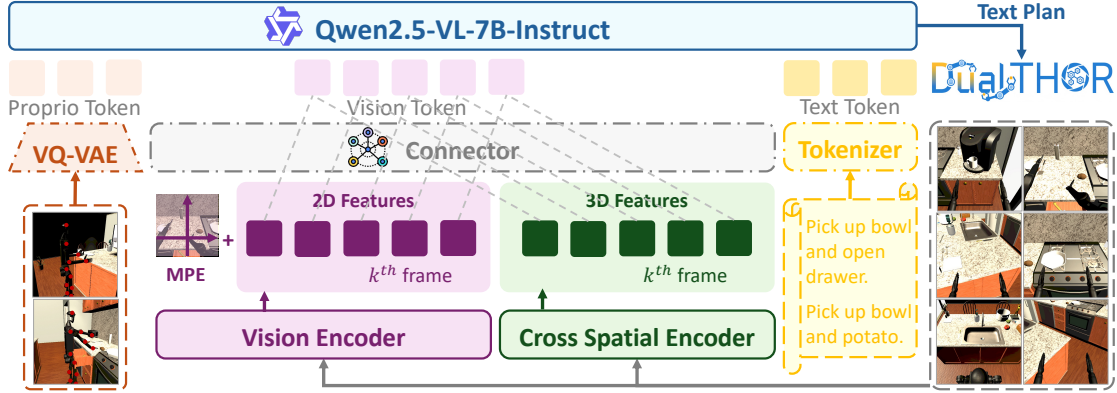


Fig. 4: **The framework of Proprio-MLLM.** By incorporating proprioceptive information, we propose a multimodal alignment large language model, **Proprio-MLLM**. We introduce a motion-based position embedding method and a cross-spatial encoder, increasing the model’s embodiment awareness and spatial reasoning in dual-arm tasks.

loss. The overall loss is formulated as:

$$\mathcal{L} = \|D(z(m)) - m\|^2 + \alpha \|sg[Z(m)] - e\|_2^2 + \beta \|Z(m) - sg[e]\|_2^2, \quad (2)$$

where $sg[\cdot]$ denotes the stop-gradient operation and α, β are hyper-parameters to control the relative weight of different components. The codebook is updated using an Exponential Moving Average (EMA) strategy to improve stability.

Stage 2: Alignment Training. We combine all modalities into a unified token space, utilizing shared embedding layers and attention mechanisms. The attention mechanisms of concatenated hidden states $H_{m,i,t} = [H_m, H_i, H_t]$ including motion, image and text embeddings are calculated by,

$$\begin{aligned} Q_{m,i,t} &= W_Q H_{m,i,t}, & K_{m,i,t} &= W_K H_{m,i,t}, \\ V_{m,i,t} &= W_V H_{m,i,t}. \end{aligned} \quad (3)$$

where $\{Q, K, V\}_{m,i,t}$ represents query, key and value, $W_{\{Q,K,V\}}$ represents their weight matrices.

Since MLLMs already exhibit strong image-text alignment, we initially freeze the image and text weights separately to improve motion alignment of single modality to prevent treating motion as noise. After training stabilizes, we perform joint training across all modalities.

Stage 3: Instruction Tuning. Finally, we use a self-learning method SELU [31] to collect some dual-arm planning data, therefore grounding the MLLM into dual-arm planning tasks with the understanding of all three modalities.

C. Motion-Based Position Embedding

To better integrate embodiment information and strengthen the correlation between motion and visual inputs, thereby enhancing Proprio-MLLM’s long-horizon reasoning based on the robot’s body state, we introduce a motion-based position embedding (MPE) for visual

tokens. Prior methods such as M-RoPE [27] encode temporal, height, and width positions, where temporal IDs remain constant for image and spatial IDs are assigned by pixel location. However, this design is less effective for embodied tasks, as it ignores interactive regions around the robot and their relative spatial relationships.

To overcome this limitation, we replace the original dimensions (temporal, height, width) with (temporal, robot-centric height, robot-centric width), where spatial indices are encoded radially outward from the robot’s centroid. Specifically, for a visual token at (x, y) and a robot centroid at (x_r, y_r) , we define:

$$MPE(t, x, y) = [t, \text{sign}(y - y_r), \text{sign}(x - x_r)], \quad (4)$$


where t denotes the temporal index. The MPE is added to visual tokens, enabling the model to distinguish tokens directions relative to the robot.

D. Cross-Spatial Encoder


While proprioceptive integration and motion-based position embedding allow Proprio-MLLM to better capture dual-arm selection logic and adapt planning to the current embodiment, a limitation remains in estimating the robot’s interactive range. Although the model can identify task-relevant objects and adjust posture, it still struggles to determine whether an object is within feasible reach.

To address this, we introduce a cross-spatial encoder (CSE) to enhance interactive range prediction. CUT3R [32] provides an accurate single-frame depth estimation, yielding reliable geometric cues when combined with arm-length constraints. Specifically, we retain the Qwen2.5-VL image encoder for high-level visual features and augment it with 3D point-map features predicted by CUT3R from a single RGB image. Given an observation $o_i \in \mathbb{R}^{H \times W \times 3}$, we obtain:


$$F_{2D} = E_{\text{Qwen2.5-VL}}(o_i), \quad F_{3D} = E_{\text{CUT3R}}(o_i), \quad (5)$$



(a) Distribution of interactive object types in DualTHOR.



(b) Distribution of single-arm interaction tasks in DualTHOR.



(c) Distribution of the three task categories in DualTHOR.

Fig. 5: Distribution of objects and tasks.

TABLE II: Performance of baselines across different task configurations in DualTHOR. Each task includes results for both X1 and H1 robots. The evaluation metric is the success rate over 50 trials for each task.

	Dual-Arm Essential Tasks		Dual-Arm Optional Tasks		Single-Arm Tasks	
	X1	H1	X1	H1	X1	H1
GPT-4o	23.31%	27.07%	39.76%	40.96%	51.67%	56.67%
Gemini-1.5-Pro	25.56%	26.33%	37.35%	36.14%	41.67%	43.33%
Qwen2.5-VL-7B-Ins	18.05%	17.29%	21.08%	19.28%	23.33%	25.00%
InternVL2.5-8B	9.77%	15.79%	23.49%	25.30%	23.33%	31.67%
LLM-Planner	28.57%	31.58%	43.37%	45.78%	55.00%	56.67%
RAP	27.81%	33.51%	45.18%	47.59%	53.33%	55.00%
DAG-Plan	36.09%	41.53%	51.20%	52.41%	55.00%	58.33%
Proprio-MLLM	59.39%	63.16%	71.69%	70.48%	73.33%	75.00%

where F_{2D} denotes 2D visual features and F_{3D} represents the point-map based features. The 3D features are reshaped, aligned with 2D features, and fused via a lightweight MLP:

$$T_{\text{vision}} = \text{MLP}[(F_{2D} + \Phi(F_{3D}))], \quad (6)$$

where $\Phi(\cdot)$ reshapes 3D features into the 2D-compatible space. The final fused vision token T_{vision} , combined with proprioceptive states, provides enhanced awareness of the robot’s interactive range, enabling more accurate dual-arm interaction planning.

V. EXPERIMENTS

A. Experimental Setup

We conduct experiments on 359 tasks across 10 distinct room environments, involving 68 unique objects. Distributions of object categories, task categories, and the proportion of each task type are illustrated in Fig. 5. Most dual-arm essential tasks and dual-arm optional tasks are composed of multiple single-arm tasks, while a smaller portion is manually designed. We run all experiments in 4 x A100 GPUs with 40GB memory, and it takes about 8 hours to get the results.

B. Baselines

We evaluate three distinct classes of baselines in DualTHOR: proprietary MLLMs, open-source MLLMs, and prompt-enhanced MLLMs. The planning capabilities of these models for dual-arm tasks are assessed mainly by their success rates across all task categories.

Proprietary MLLMs (GPT-4o [33], Gemini 1.5 Pro [34]) are closed-source commercial systems trained on large-scale private datasets and optimized for downstream tasks. Open-source MLLMs (Qwen2.5-VL-7B [27], InternVL2.5-8B [35]) are publicly available models, typically trained on transparent datasets. Prompt-enhanced MLLMs (LLM Planner [1], RAP [15], DAG-Plan [4]) leverage structured prompts to improve reasoning in embodied planning. Base model used is GPT-4o. DAG-Plan [4] is the state-of-the-art method for dual-arm planning that utilizes a graph-based structure to represent dual-arm skills.

C. Main Results for Dual-Arm Tasks

We test all baselines across the three task categories defined in Section III-B. Each task is tested on both X1 and H1 robots with 50 trials to ensure statistical reliability. Initial positions vary across rooms but remain consistent within the same room to preserve the reliability of historical trajectories for prompt-enhanced MLLMs. All baselines are tested under identical conditions: temperature set to 0, maximum token length 2048, input images scaled to 500×500 resolution, and a maximum of 50 environment steps for high-level planning.

Is proprioceptive information necessary for long-horizon dual-arm planning tasks? Our results in Table II show that existing methods struggle with long-horizon dual-arm planning. By incorporating proprioceptive information along with the proposed motion-based position embedding and cross-spatial encoder,

Proprio-MLLM achieves the best overall performance. Compared to DAG-Plan, which leverages graph structures to represent robot skills, Proprio-MLLM improves average performance by 19.75%, and up to 23.30% on dual-arm essential tasks for X1 robot. Across the three task categories, we observe a pronounced gap between dual-arm and single-arm planning in current MLLMs, highlighting the scarcity of dual-arm humanoid planning data in existing simulators and motivating the development of DualTHOR for more diverse embodied planning data. Notably, proprioceptive information consistently enhances performance across all tasks, with the largest gains observed in dual-arm scenarios.




Fig. 6: **Impact of proprioceptive integration.** We analyze the failed trajectories of the basic model, the state-of-the-art method, and our Proprio-MLLM, categorizing these failures into three types to assess how each model leverages proprioceptive information.

How does proprioceptive information influence dual-arm high-level planning? To highlight the impact of proprioceptive information on high-level planning, we analyze failure trajectories and categorize them into three types: navigation failure (unable to reach the target interaction point), body adjustment failure (mismatched height with the target object leading to low-level execution failure), and logical failure (suboptimal arm selection in long-horizon planning). Statistics in Fig. 6 show that both DAG-Plan and Proprio-MLLM reduce failures, indicating that proprioceptive information improves arm selection logic. However, DAG-Plan’s graph-based prompting cannot fully exploit proprioception for other failure types. This limitation motivates the use of motion-based position embedding and the cross-spatial encoder, which together enable Proprio-MLLM to integrate image perception and proprioceptive states more effectively, producing coherent and adaptive dual-arm plans.

How effectively can MLLMs handle contingency? Since DualTHOR provides both continuous interaction and a contingency mechanism, our benchmark can evaluate the re-planning ability of MLLMs, which is fundamental to ensuring the safety of high-level planning in real-world deployment. Each task is further divided

TABLE III: **Performance of re-planning in DualTHOR.** Tasks are categorized into three difficulty levels based on the success rate of low-level skills.

Method	H1		
	Easy	Medium	Hard
Proprio-MLLM	63.16%	51.63%	36.51%
Proprio (w/o reflection)	56.39%	38.67%	26.17%
DAG-Plan	41.53%	20.33%	15.53%
DAG (w/o reflection)	34.81%	10.97%	7.26%

into three difficulty levels — Easy (100% action success rate), Medium (50%), and Hard (20%). To accomplish such tasks, agents must analyze failure scenarios and re-plan accordingly. For example, they should locate a new cup when the original one is broken. We evaluate the performance of DAG-Plan and Proprio-MLLM, with and without the reflection prompt, across three categories, and report the representative results on H1 dual-arm essential tasks in Table III. From the experimental results, we observe that Proprio-MLLM leverages proprioceptive information more effectively, and the reflection prompt contributes more substantially to improving re-planning ability. However, the limited performance of both methods on high-difficulty tasks reveals the scarcity of contingency cases in the training data of current MLLMs. This limitation further underscores the importance of DualTHOR in providing a contingency mechanism for the development of robust planning datasets.

TABLE IV: **Results of ablation study of different components in Proprio-MLLM.** We examine the impact of MPE, CSE, and proprioceptive integration by analyzing the reduction of three failure categories.

Method	Success Rate	Failure Categories		
		#1	#2	#3
Proprio-MLLM	63.16%	1225	784	441
Proprio w/o CSE	45.17%	2205	902	538
Proprio w/o MPE	52.33%	1585	951	634
Proprio w/o ALL	41.53%	2238	972	678
Qwen2.5-VL-Ins	17.29%	2751	1650	1100

D. Ablation Study

To investigate the impact of each component in **Proprio-MLLM**, we conduct ablation experiments by removing the motion-based position embedding (Proprio w/o MPE) and cross-spatial encoder (Proprio w/o CSE), as shown in Table IV. Proprio w/o ALL refers to removing both components, isolating the effect of proprioceptive information integration. Failure #1–3 correspond to navigation, body adjustment, and logical failures. The results show that incorporating proprioceptive information significantly enhances MLLM planning by aligning with the robot’s body state. Specifically, Proprio w/o CSE emphasizes the cross-spatial encoder’s role in improving the robot’s interaction range, reducing navigation failures

(orange squares). In contrast, Proprio w/o MPE highlights the importance of motion-based position embedding in adjusting the robot’s body, minimizing arm selection logical errors and reducing the latter two failure types (blue squares). Combining both components enables Proprio-MLLM to achieve proprioception-aware, long-horizon planning for dual-arm humanoid robots.

VI. CONCLUSION

We present DualTHOR, a high-fidelity planning simulation environment designed to enable continuous and realistic bimanual humanoid robot interaction. We design three categories of dual-arm tasks to evaluate planning ability of MLLMs. To optimize the dual-arm planning performance, we propose Proprio-MLLM to utilize the proprioceptive information into long-horizon plan. We introduce motion-based position embedding and cross-spatial encoder to understand the dual-arm humanoid body state and interaction range to achieve proprioception-aware plan and reduce three kinds of failures. Proprio-MLLM can achieve an average improvement of 19.75% in planning performance, demonstrating more reliable logical and spatial reasoning capabilities.

REFERENCES

- [1] C. H. Song, J. Wu, C. Washington, *et al.*, “Llm-planner: Few-shot grounded planning for embodied agents with large language models,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 2998–3009.
- [2] Z. Fu, Q. Zhao, Q. Wu, *et al.*, “Humanplus: Humanoid shadowing and imitation from humans,” in *Conference on Robot Learning*, 2024.
- [3] Y. Chen, W. Cui, Y. Chen, *et al.*, “Robogpt : An intelligent agent of making long-horizon decisions for instruction daily tasks,” *IEEE Transactions on Cognitive and Developmental Systems*, 2025.
- [4] Z. Gao, Y. Mu, J. Qu, *et al.*, “Dag-plan: Generating directed acyclic dependency graphs for dual-arm cooperative planning,” *arXiv preprint arXiv:2406.09953*, 2024.
- [5] Z. Mandi, S. Jain, and S. Song, “Roco: Dialectic multi-robot collaboration with large language models,” in *2024 IEEE International Conference on Robotics and Automation*, 2024.
- [6] M. Deitke, W. Han, A. Herrasti, *et al.*, “Robothor: An open simulation-to-real embodied ai platform,” *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [7] M. Deitke, E. VanderBilt, A. Herrasti, *et al.*, “Proctor: Large-scale embodied ai using procedural generation,” *Advances in Neural Information Processing Systems*, 2022.
- [8] Y. Chen, S. Tian, S. Liu, *et al.*, “Conrft: A reinforced fine-tuning method for vla models via consistency policy,” in *Proceedings of Robotics: Science and Systems*, 2025.
- [9] E. Kolve, R. Mottaghi, W. Han, *et al.*, “Ai2-thor: An interactive 3d environment for visual ai,” *arXiv preprint arXiv:1712.05474*, 2017.
- [10] K. Ehsani, W. Han, A. Herrasti, *et al.*, “Manipulathor: A framework for visual object manipulation,” *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [11] P. Zhang, X. Gao, Y. Wu, *et al.*, “Moma-kitchen: A 100k+ benchmark for affordance-grounded last-mile navigation in mobile manipulation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 6315–6326.
- [12] C. Sferrazza, D.-M. Huang, X. Lin, *et al.*, “Humanoidbench: Simulated humanoid benchmark for whole-body locomotion and manipulation,” *arXiv preprint arXiv:2403.10506*, 2024.
- [13] V. Makoviychuk, L. Wawrzyniak, Y. Guo, *et al.*, “Isaac gym: High performance gpu-based physics simulation for robot learning,” *arXiv preprint arXiv:2108.10470*, 2021.
- [14] H. Li, Z. Jiang, Y. Chen, and D. Zhao, “Generalizing consistency policy to visual rl with prioritized proximal experience regularization,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 109672–109700, 2024.
- [15] T. Kagaya, T. J. Yuan, Y. Lou, *et al.*, “Rap: Retrieval-augmented planning with contextual memory for multimodal llm agents,” in *NeurIPS Workshop on Open-World Agents*, 2024.
- [16] Z. Luo, J. Cao, K. Kitani, *et al.*, “Perpetual humanoid control for real-time simulated avatars,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [17] R. Gong, X. Gao, Q. Gao, *et al.*, “Lemma: Learning language-conditioned multi-robot manipulation,” *IEEE Robotics and Automation Letters*, 2023.
- [18] B. Li, H. Li, Y. Zhu, and D. Zhao, “Mat: Morphological adaptive transformer for universal morphology policy learning,” *IEEE Transactions on Cognitive and Developmental Systems*, vol. 16, no. 4, pp. 1611–1621, 2024.
- [19] C. Gan, J. Schwartz, S. Alter, *et al.*, “Threedworld: A platform for interactive multi-modal physical simulation,” *arXiv preprint arXiv:2007.04954*, 2020.
- [20] C. Li, F. Xia, R. Martín-Martín, *et al.*, “igibson 2.0: Object-centric simulation for robot learning of everyday household tasks,” *arXiv preprint arXiv:2108.03272*, 2021.
- [21] C. Li, R. Zhang, J. Wong, *et al.*, “Behavior-1k: A benchmark for embodied ai with 1,000 everyday activities and realistic simulation,” in *Conference on Robot Learning*, 2023, pp. 80–93.
- [22] J.-B. Alayrac, J. Donahue, P. Luc, *et al.*, “Flamingo: a visual language model for few-shot learning,” *Advances in neural information processing systems*, 2022.
- [23] W. Dai, N. Lee, B. Wang, *et al.*, “Nvllm: Open frontier-class multimodal llms,” *arXiv preprint arXiv:2409.11402*, 2024.
- [24] H. Liu, C. Li, Q. Wu, *et al.*, “Visual instruction tuning,” *Advances in neural information processing systems*, 2023.
- [25] B. Jiang, X. Chen, W. Liu, *et al.*, “Motiongpt: Human motion as a foreign language,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 20067–20079, 2023.
- [26] M. Pan, J. Zhang, T. Wu, *et al.*, “Omnimanip: Towards general robotic manipulation via object-centric interaction primitives as spatial constraints,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025.
- [27] S. Bai, K. Chen, X. Liu, *et al.*, “Qwen2.5-vl technical report,” *arXiv preprint arXiv:2502.13923*, 2025.
- [28] C. Guo, S. Zou, X. Zuo, *et al.*, “Generating diverse and natural 3d human motions from text,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 5152–5161.
- [29] M. Loper, N. Mahmood, J. Romero, *et al.*, “Smpl: A skinned multi-person linear model,” in *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 2023, pp. 851–866.
- [30] A. Van Den Oord, O. Vinyals, *et al.*, “Neural discrete representation learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [31] B. Li, H. Jiang, Z. Ding, *et al.*, “SELU: Self-learning embodied multimodal large language models in unknown environments,” *Transactions on Machine Learning Research*, 2025.
- [32] Q. Wang, Y. Zhang, A. Holynski, *et al.*, “Continuous 3d perception model with persistent state,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 10510–10522.
- [33] A. Hurst, A. Lerer, A. P. Goucher, *et al.*, “Gpt-4o system card,” *arXiv preprint arXiv:2410.21276*, 2024.
- [34] G. Team, R. Anil, S. Borgeaud, *et al.*, “Gemini: a family of highly capable multimodal models,” *arXiv preprint arXiv:2312.11805*, 2023.
- [35] Z. Chen, W. Wang, Y. Cao, *et al.*, “Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling,” *arXiv preprint arXiv:2412.05271*, 2024.