

TIGeR: Tool-Integrated Geometric Reasoning in Vision-Language Models for Robotics

Yi Han^{1,2,*}, Enshen Zhou^{1,2,*†}, Shanyu Rong^{2,3}, Jingkun An¹,
Pengwei Wang², Zhongyuan Wang², Cheng Chi^{2,†}, Lu Sheng^{1,2,†}, Shanghang Zhang^{2,3,†}

Abstract—Vision-Language Models (VLMs) have shown remarkable capabilities in spatial reasoning, yet they remain fundamentally limited to qualitative precision and lack the computational precision required for real-world robotics. Current approaches fail to leverage metric cues from depth sensors and camera calibration, instead reducing geometric problems to pattern recognition tasks that cannot deliver the centimeter-level accuracy essential for robotic manipulation. We present *TIGeR* (Tool-Integrated Geometric Reasoning), a novel framework that transforms VLMs from perceptual estimators to geometric computers by enabling them to generate and execute precise geometric computations through external tools. Rather than attempting to internalize complex geometric operations within neural networks, *TIGeR* empowers models to recognize geometric reasoning requirements, synthesize appropriate computational code, and invoke specialized libraries for exact calculations. To support this paradigm, we introduce *TIGeR-300K*, a comprehensive tool-invocation-oriented dataset covering point transformations, pose estimation, trajectory generation, and spatial compatibility verification, complete with tool invocation sequences and intermediate computations. Through a two-stage training pipeline combining supervised fine-tuning (SFT) and reinforcement fine-tuning (RFT) with our proposed hierarchical reward design, *TIGeR* achieves SOTA performance on geometric reasoning benchmarks while demonstrating centimeter-level precision in real-world robotic manipulation tasks. See the project page at <https://hany01rye.github.io/TIGeR/>.

I. INTRODUCTION

Robots operating in the physical world require *geometric reasoning*—the ability to compute metric quantities such as 3D poses, rotation matrices, and collision-free trajectories with centimeter precision. Yet despite recent progress, current Vision-Language Models (VLMs) remain limited to spatial reasoning [1], [2], [3], providing only qualitative assessments of relations such as “left of” or “reachable”, while lacking the computational machinery for quantitative accuracy. This limitation prevents embodied VLMs from supporting the exact manipulations and planning that real-world robotics demands.

This limitation manifests on both the perception and output modalities. On the perception side, although depth sensors and stereo cameras provide rich metric information, existing models typically reduce these signals to image-like representations [4], [5], [6], [7], discarding the geometric properties required for precise computation. Without explicit

use of camera intrinsics and extrinsics, they cannot transform pixel observations into real-world coordinates, leaving them incapable of executing metric operations such as pose estimation or point cloud alignment. On the output side, the problem is even more pronounced. A few recent efforts attempt to predict 3D bounding boxes or poses [8], yet these results are statistical regressions trained from data rather than deterministic computations grounded in geometry. The majority of VLMs remain constrained to qualitative descriptions or 2D pixel predictions, which are insufficient for generating the precise poses, trajectories, and kinematic constraints that robotic manipulation and motion planning demand. *In essence, existing approaches collapse geometric reasoning into pattern recognition, sacrificing the computational fidelity required for real-world robotics.*

To overcome this limitation, we introduce *TIGeR* (Tool-Integrated Geometric Reasoning), a framework that re-defines the role of VLMs from perceptual estimators to geometric computers. Rather than attempting to internalize geometric calculations within neural networks, *TIGeR* enables models to detect when geometric reasoning is required, generate code for the appropriate operations, and invoke external tools to execute them. Through integration with depth sensors, camera intrinsics, and geometric libraries, VLMs can perform exact computations that significantly go beyond approximation.

To support this paradigm, we construct *TIGeR-300K*, a dataset of 300K samples spanning point transformations, pose estimation, trajectory generation, and spatial compatibility checks. Each sample provides not only problem statements and solutions but also the complete tool invocation sequence and intermediate computations. Building on *TIGeR-300K*, we develop a two-stage training pipeline: supervised fine-tuning (SFT) to instill tool usage, followed by reinforcement fine-tuning (RFT) with our proposed hierarchical reward to enhance the accuracy and task completion.

In summary, this work makes three contributions:

- *Concept & Method.* We underscore the central role of *geometric reasoning*—beyond qualitative spatial reasoning—for precise robotic control, and introduce *TIGeR*, a tool-integrated framework that enables exact geometric computation via code generation and execution on calibrated metric inputs.
- *Dataset.* We release *TIGeR-300K*, a large-scale, tool-invocation-oriented dataset explicitly designed for geometric reasoning via programmatic tool calls.
- *Experimental Results.* Using a two-stage SFT→RFT

*These authors contributed equally. †Project Leader.

‡Corresponding author: chicheng@baai.ac.cn,
lsheng@buaa.edu.cn, shanghang@pku.edu.cn

¹ Beihang University ² Beijing Academy of Artificial Intelligence ³ State Key Laboratory of Multimedia Information Processing, School of Computer Science, Peking University

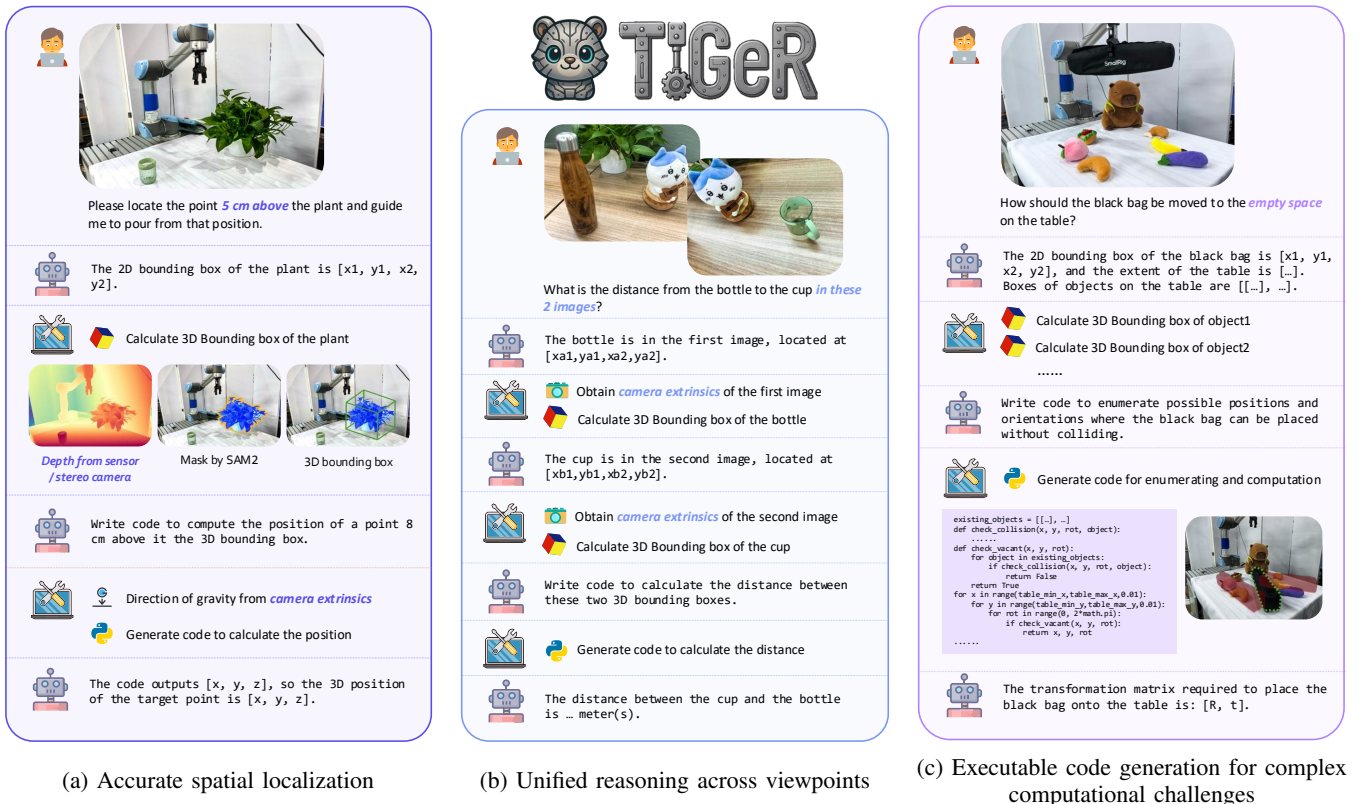


Fig. 1: *TIGeR* is a tool-integrated framework that enables exact geometric reasoning and computation via code generation and execution on calibrated metric inputs. It can achieve accurate spatial localization, unified reasoning across viewpoints, and address complex embodied tasks requiring accurate numerical computation, with high interpretability and adaptability.

pipeline, *TIGeR* achieves state-of-the-art performance on geometric benchmarks and achieves centimeter-level accuracy in real-world robotic manipulation.

II. RELATED WORK

Tool-Integrated Reasoning. Tool-Integrated Reasoning (TIR) enhances LLM and VLM reasoning ability, particularly for tasks requiring precise numerical computation, and is categorized into prompt-based, SFT-based, and RL-based methods. Prompt-based methods [9], [10], [11] use carefully designed prompts to invoke external tools without additional training. While easy to implement, they often exhibit low accuracy and unstable tool usage. SFT-based methods [12], [13], [14], [15], [16], [17] fine-tune pre-trained models on TIR-specific data, improving performance but risking memorizing tool usage patterns from training data, which may hinder generalization to unseen scenarios. While current RL-based methods [18], [19], [20] promote generalization through outcome-based rewards, they overlook intermediate reasoning precision, essential for embodied tasks involving accurate numerical computation. We thus propose novel process-based reward functions tailored for TIR and introduce a new dataset to support training.

Spatial Understanding and Reasoning with VLMs. Spatial understanding focuses on object-centric properties (*e.g.*, position, orientation) and inter-object relations (*e.g.*, distance, direction), while spatial reasoning draws higher-level infer-

ences over such information. Recent advances in VLMs enhance these two abilities via two paradigms: **(1)** Data-driven methods that fine-tune VLMs with pseudo-3D annotations [21], [4], real-world 3D datasets [22], or simulated data [23], using either the original image encoder, an added depth encoder [4], [7], [6] for 3D cues, or a geometry encoder [2] for structured scene and camera information. While these methods may enhance spatial understanding, they struggle with embodied tasks requiring precise numerical reasoning due to VLM’s hallucination and limited interpretability. **(2)** Tool-based methods [24] that integrate VLMs with vision foundation models to extract and reason spatial cues can achieve better performance but suffer from a limited training-free approach. In contrast, this work introduces a sequential SFT-RFT training pipeline with novel process reward functions, to enable spatial understanding and reasoning with accurate numerical computation.

III. METHOD

We first give an overview of the proposed *TIGeR* (Sec. III-A). Next, we introduce the tool categorization in embodied scenarios (Sec. III-B). Then, we describe the construction of *TIGeR-300K* (Sec. III-C). Finally, we elaborate on the two-stage training pipeline (Sec. III-D), including our novel hierarchical reward design tailored for TIR.

A. Overview of TIGeR

TIGeR is a tool-integrated geometric reasoning framework that performs exact geometric computation by generating and executing code on calibrated metric inputs, enabling high-precision reasoning for downstream robotic tasks, even at centimeter-level accuracy. Specifically, as shown in Fig 1(c), given a human instruction (e.g., “How should the black bag be moved to the empty space on the table?”), TIGeR performs step-by-step reasoning by selecting relevant context (e.g., previous tool outputs) and invoking appropriate tools (e.g., 2D/3D detection, camera pose estimation) to obtain desired precise intermediate metric results. During reasoning, it can also generate code to integrate and compute over acquired information, enabling the progressive derivation of informative intermediate results and final answers.

In embodied robotics tasks demanding precise geometric reasoning, our paradigm offers three key advantages over directly using data-driven, fine-tuned, spatially-aware VLMs: **(1) Accurate spatial localization.** Our method enables precise prediction of positions in 3D space (e.g., 3D key-points, trajectories) by leveraging existing 3D foundation models. In contrast, data-driven VLMs often struggle with quantitative spatial reasoning due to training objectives (i.e., next-token prediction) and insufficient data coverage (e.g., understanding “5cm above the plant”). **(2) Unified reasoning across viewpoints.** Our method is viewpoint-invariant and supports consistent numerical reasoning within a unified coordinate system, even when multi-view cameras are not jointly calibrated, enabled by recent advances in 3D geometry reconstruction models (e.g., π^3 [25]). **(3) Interpretability and Adaptability.** By explicitly invoking tools and exposing intermediate steps, our method offers explicit reasoning and easy integration of new tools (e.g., SOTA models). This enables efficient adaptation to novel tasks without costly retraining, unlike monolithic data-driven models requiring retraining to generalize to new scenarios.

B. Tool Categorization in Embodied Scenarios

We categorize tools into two groups: visual perception and geometric computation, reflecting a hierarchical workflow. Visual perception tools extract pixel- or camera-level information from sensory inputs, while geometric computation tools process this data to derive geometric properties (e.g., point transformations, pose estimation). The computed results are then translated into executable robot commands. This organization enables efficient tool invocation, where the VLM determines when and how to invoke tools based on task demands, explicitly injecting precise geometric information into the reasoning process.

Visual Perception Tools extract essential sensory information to address the limitations of VLMs in metric-scale accuracy. Notably, in real-world settings where certain data (e.g., camera intrinsics and calibrated extrinsics) are readily available, we directly use these values instead of invoking additional foundation models as tools. Key tools include:

- `camera_intrinsics`: Retrieves the camera’s intrinsic parameters to support accurate 2D-to-3D projections. If

unavailable, we estimate it using MoGe-2 [26].

- `camera_extrinsics`: Provides extrinsic parameters to align multi-view observations within a unified 3D coordinate system. When not directly accessible, we estimate them using VGGT [27] or π^3 [25], referencing the first frame as the world coordinate system.
- `depth_sensor`: Queries depth information from sensors for selected regions, avoiding the overhead of inputting full-depth maps to the VLM. If sensor data is missing, metric depth is estimated via MoGe-2 [26].
- `object_segmentation`: Generates precise segmentation masks from coarse VLM outputs (e.g., 2D bounding boxes), enabling fine-grained analysis of object shapes and boundaries. We use SAM2 [28] as the tool.

Geometric Computation Tools perform geometric computations and transformations, bridging the gap between 2D pixel-based VLM outputs and 3D geometric requirements. These tools can handle embodied tasks requiring accurate numerical computation (e.g., 3D pose estimation) beyond the VLM’s native capabilities. Notably, to enhance accuracy and robustness, we introduce a code generation subroutine: the VLM specifies high-level requirements, which are fulfilled by an external code generator (Qwen3-Coder [29]). The generated code, executed in a sandboxed environment, supports progressive reasoning by integrating intermediate results and producing final outputs. Key tools include:

- `box_2d_to_box_3d`: Converts a 2D bounding box (from VLM output) into a 3D bounding box by integrating 2D segmentation with depth sensor data, estimating object scale, position, and orientation in 3D space.
- `point_3d_to_point_2d`: Projects 3D points onto the 2D image plane using camera intrinsics/extrinsics, enabling tasks such as pixel-level target location referring.
- `code_executor`: Invokes the code generation pipeline for arbitrary computations, such as inter-object distance calculations or pose estimations, returning numerical results to the VLM’s reasoning process.

C. Data Generation

To support effective training of TIGeR, we construct a large-scale dataset, TIGeR-300K, comprising 300K high-quality samples that integrate tool usage with geometric reasoning in embodied scenarios. As shown in Fig. 2, our data generation pipeline adopts a hybrid approach: **(1)** a structured, template-based annotation method for generating precise and controllable instances targeting common spatial queries, and **(2)** a large-model-driven chain-of-thought rewriting process to produce diverse and flexible examples that promote adaptive tool-integrated reasoning. Each sample provides not only problem statements and solutions but also the complete tool invocation sequence and intermediate computations. This dual strategy ensures comprehensive coverage of core geometric concepts while supporting generalization to diverse, real-world spatial and geometric reasoning tasks. **Template-Based Data Generation.** We utilize the CA-1M dataset [30] as our primary data source, which provides

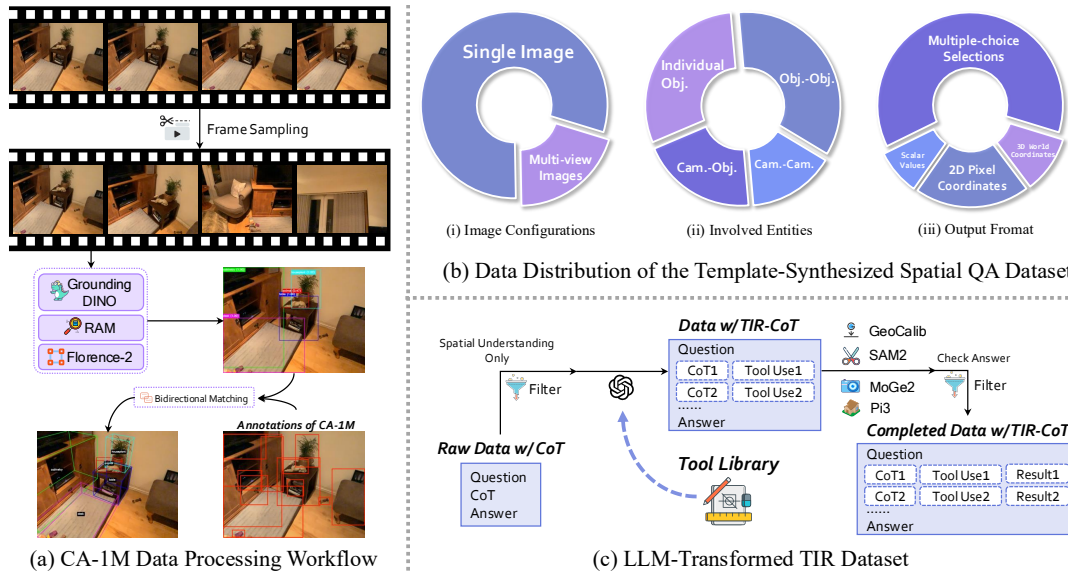


Fig. 2: *TIGeR-300K*: A VQA dataset comprising 300K samples with Tool-Integrated Geometric Reasoning, generated via template-based synthesis and tool-augmentation by LLMs. It contains problem statements and solutions with a complete tool invocation sequence and intermediate computations tailored for our proposed hierarchical reward design.

per-frame 2D/3D oriented bounding boxes, camera intrinsics/extrinsics, and depth maps. To reduce redundancy, we sample every 20th frame. However, the original dataset presents several challenges, including numerous bounding boxes with ambiguous or semantically irrelevant content, as well as a lack of semantic labels. To address these issues, we perform data cleaning and re-annotation. Specifically, as shown in Fig 2(a), we leverage GroundingDINO [31], RAM [32], and Florence-2 [33] to predict semantic labels and 2D bounding boxes for previously unlabeled objects. The predicted boxes are then matched with the original CA-1M annotations via Intersection over Union (IoU), retaining only those with high overlap to filter out noisy boxes and enrich the dataset with semantic information. Finally, we construct 3D scene graphs—where nodes represent objects and edges encode inter-object spatial relations—by integrating cleaned 2D/3D boxes, depth maps, and camera parameters. These graphs serve as a foundation for downstream template-based data generation. To address common spatial queries in robotics (*e.g.*, localization, distance estimation, and pose computation), we design modular templates that vary systematically to generate diverse yet consistent instances, filled using annotated CA-1M frames. These templates exhibit diversity (see Fig 2(b)) in the following aspects:

- **Image Configurations:** Samples comprise both single-view and multi-view images from the same scene, enabling diverse embodied task settings.
- **Involved Entities:** Queries involve object-centric properties (*e.g.*, size), inter-object relations (*e.g.*, distance, spatial layout), object-camera interactions (*e.g.*, depth), and inter-camera geometry (*e.g.*, relative poses).
- **Output Formats:** Answers are structured to support both robotic applications (*e.g.*, 3D coordinates, poses) and spatial reasoning benchmarks (*e.g.*, multiple choice, scalar values, 2D pixel locations).

All templates are instantiated using filtered and annotated CA-1M frames, ensuring access to ground-truth geometry for accurate tool simulation. This yields approximately 274K structured samples, forming a robust dataset for supervised learning of geometric reasoning tasks.

Large-Model-Based CoT Rewriting. To improve adaptability and cover a wider range of geometric reasoning tasks, we utilize the SSR-CoT dataset [5], a visual question answering (VQA) collection enriched with chain-of-thought (CoT) annotations. As shown in Fig. 2(c), we first apply GPT-4o to filter for spatial inference-related questions (*e.g.*, depth estimation, object relations, and multi-view consistency), yielding a focused subset. For each selected CoT instance, we prompt a large model to rewrite it into a TIR-enhanced version, inserting tool calls where geometric precision is needed. Tool responses are initially left as placeholders. Since SSR-CoT lacks ground-truth annotations, we invoke tools by constructing task-relevant functions using foundation models like MoGe2 [26] and GeoCalib [26], and populate the placeholders accordingly. Quality checks are performed to ensure valid tool formats and accurate final answers, resulting in 35K diverse, tool-integrated samples. By combining template-generated and rewritten data, *TIGeR-300K* provides the VLM with both precise supervision and compositional flexibility, enhancing its capacity to parse unstructured queries and adapt tool usage in novel scenarios.

D. Two-Stage Training

We build on GLM-4.1V-Thinking [34] and introduce tool-integrated geometric reasoning for robotics via a two-stage training pipeline (Fig. 3): Supervised Fine-Tuning (SFT) followed by Reinforcement Fine-Tuning (RFT). Leveraging *TIGeR-300K*, SFT imparts basic tool-use reasoning capabilities, while RFT refines them through reward signals focused on geometric computation accuracy and effective tool use.

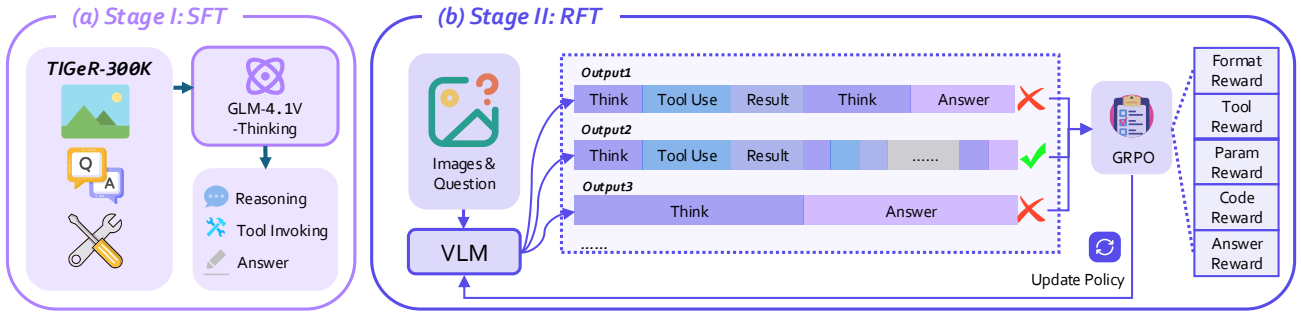


Fig. 3: *TIGeR*'s two-stage training pipeline. The process begins with SFT on *TIGeR-300K*, followed by RFT incorporating five specialized reward functions (outcome- and process-based) tailored for TIR to enhance geometric reasoning abilities.

Cold-Start SFT. We first apply SFT on *TIGeR-300K* to initialize the model's ability to generate reasoning chains with appropriate tool usage. Each sample in *TIGeR-300K* has four elements: input image(s) \mathcal{I} , textual query \mathcal{Q} , reasoning trajectory \mathcal{R} (including tool calls and code generation), and final output \mathcal{O} (e.g., computed geometric results). We optimize the model using a next-token prediction loss:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{\tau \sim \mathcal{D}} \sum_{t=1}^T \log \pi_{\theta}(r_t \mid \mathbf{I}, \mathbf{Q}, \mathbf{r}_{<t})$$

where \mathcal{D} is *TIGeR-300K*, π_{θ} denotes the parameterized VLM, r_t is the token at position t in the reasoning-output sequence of length T , and $\mathbf{r}_{<t}$ are the preceding tokens.

RFT for Geometric Tool Use. We further employ RFT using GRPO [35] to enhance the model's generalized tool-integrated geometric reasoning capabilities. For a given visual-language query q and a batch of N responses $\{y_i\}_{i=1}^N$ sampled from the current policy π_{θ} , we optimize:

$$\mathcal{L}_{\text{GRPO}}(\theta) = -\mathbb{E}_{q, \{y_i\} \sim \pi_{\theta_{\text{old}}}} \left[\frac{1}{N} \sum_{i=1}^N \min \left(\rho_i \hat{A}_i, \text{clip}(\rho_i, 1 - \epsilon, 1 + \epsilon) \hat{A}_i \right) \right] + \beta D_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{SFT}})$$

where $\rho_i = \frac{\pi_{\theta}(y_i|q)}{\pi_{\theta_{\text{old}}}(y_i|q)}$ is the importance ratio, $\hat{A}_i = \frac{r_i - \bar{r}}{\sigma_r}$ is the normalized advantage with reward r_i , mean \bar{r} and standard deviation σ_r computed over the batch, and β controls KL regularization strength.

Hierarchical Reward Design. We design a composite reward function $r(y) = \sum_k w_k \cdot r_k(y)$ to evaluate geometric tool use from both outcome- and process-based perspectives: (1) **Format Reward** r_{format} (outcome-based): Measures structural correctness of spatial tokens and tool syntax: $r_{\text{format}}(y) = \mathbb{I}[\text{valid spatial tokens}] \cdot \mathbb{I}[\text{valid tool syntax}]$ (2) **Tool Invocation Reward** r_{tool} (process-based): Assesses correctness of tool selection and parameter formatting: $r_{\text{tool}}(y) = \mathbb{I}[\text{tool} \in \mathcal{T}] \cdot \mathbb{I}[\text{valid parameter structure}]$ where \mathcal{T} denotes the set of available geometric tools. (3) **Parameter Content Reward** r_{param} (process-based): Evaluates parameter accuracy based on their type:

$$r_{\text{param}}(y) = \begin{cases} \exp(-\alpha \|p - p^*\|_2) & \text{if } p \text{ is continuous} \\ \mathbb{I}[p = p^*] & \text{if } p \text{ is discrete} \end{cases}$$

where p are the predicted continuous (e.g., pixels, coordinates) and discrete (e.g., view indices) parameters, p^* are ground-truth, and α is a scaling factor. (4) **Code Generation Reward** r_{code} (process-based): For code-based geometric computations: $r_{\text{code}}(y) = \lambda_{\text{exec}} \cdot \mathbb{I}[\text{code executes}] + \lambda_{\text{out}} \cdot \mathbb{I}[\text{output correct}]$ (5) **Answer Reward** r_{answer} (outcome-based): Assesses the final answer w.r.t. the ground truth:

$$r_{\text{answer}}(y) = \begin{cases} \exp(-\gamma \|a - a^*\|) & \text{for continuous outputs} \\ \mathbb{I}[a = a^*] & \text{for discrete outputs} \end{cases}$$

This hierarchical reward structure provides fine-grained supervision for tool-integrated reasoning, guiding the model to learn accurate geometry computations and appropriate tool invocation, while preserving interpretable intermediate steps. The weights $\{w_k\}$ are tuned based on validation performance to balance structural correctness and task-level accuracy.

IV. EXPERIMENTS

Our experiments aim to address the following questions: (1) Can *TIGeR* achieve precise spatial understanding and geometry reasoning across public benchmarks (Sec. IV-A)? (2) Can *TIGeR* do spatially-aware manipulation tasks in simulators (Sec. IV-B) and on real robots (Sec. IV-C), especially requiring accurate numerical computation? (3) Which design choices greatly enhance performance (Sec. IV-D)?

A. Spatial and Geometry Reasoning Benchmarks

We evaluate *TIGeR* on qualitative spatial understanding benchmarks, including CV-Bench [36], BLINK [37], RoboSpatial [38], and EmbSpatial [39], as well as the quantitative geometric reasoning benchmark Q-Spatial++ [40]. Since these benchmarks lack ground-truth geometric annotations (e.g., camera intrinsics, extrinsics, and depth), we leverage visual foundation models (see Sec. III-B) to extract such information and inject approximate geometric priors into Tool-Integrated Reasoning at inference. As shown in Tab. I, *TIGeR* achieves a zero-shot state-of-the-art accuracy of 79.30% across these benchmarks, even outperforming Gemini 2.5-Pro by 5.83%. Fig. 4 further illustrates representative examples where *TIGeR* successfully leverages tool-integrated reasoning to solve complex spatial and geometric problems. We find that our framework can not only enhance its spatial understanding, enabled by tool-use during reasoning, but

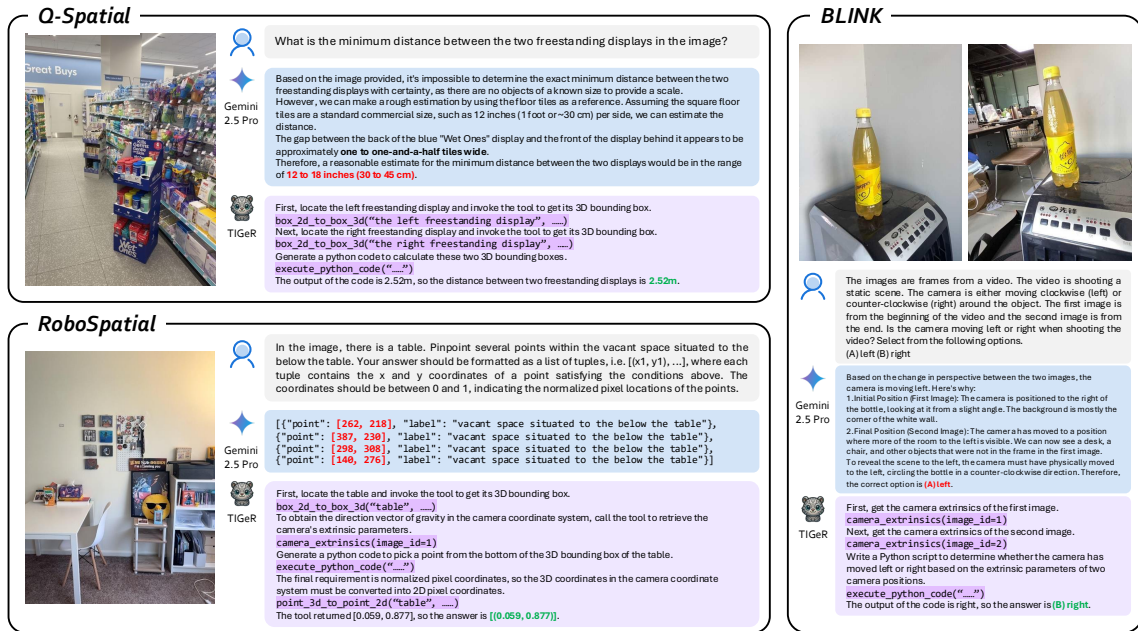


Fig. 4: Comparison between *TIGeR* and Gemini 2.5-Pro on representative samples from spatial understanding benchmarks.

TABLE I: Performance comparison on spatial understanding and geometry reasoning benchmarks across different models. M.V. indicates Multi-view parts; Conf., Comp., Cont. denote the configuration, compatibility, context part of RoboSpatial. Top-1 & Top-2 accuracies are represented using **bold text**, and underlines. For Q-Spatial++, the metric is $\delta_{\leq 2}$, where the answer is considered correct if it falls within $[0.5 \times, 2 \times]$ of the ground truth.

Models	CV-Bench [36]			BLINK [37]			RoboSpatial [38]			EmbSpatial [39]	Q-Spatial++ [40]
	2D-Rel	3D-Depth	3D-Dist	Depth	Spatial	M.V.	Conf.	Comp.	Cont.		
GPT-4o [41]	84.62	86.50	83.33	78.23	82.52	51.88	77.20	61.90	25.30	63.38	20.79
Gemini-2.5-Pro [42]	<u>93.54</u>	91.00	<u>90.67</u>	87.90	91.61	35.34	77.24	<u>68.57</u>	<u>40.20</u>	<u>76.67</u>	55.46
NVILA-8B [43]	91.54	<u>91.83</u>	<u>90.67</u>	76.61	76.92	<u>55.64</u>	59.35	23.81	5.51	67.72	50.5
Qwen-2.5-VL-7B [44]	82.15	60.17	69.00	60.98	64.34	48.12	49.59	66.67	13.61	40.20	49.50
GLM-4.1V-Thinking [34]	93.08	92.00	88.33	<u>90.32</u>	81.82	47.37	<u>81.30</u>	43.80	8.24	76.62	<u>67.33</u>
SpatialBot-3B [7]	69.38	77.33	60.83	67.74	67.83	44.36	72.36	22.86	4.88	50.66	11.88
SpaceLLaVA-13B [45]	63.69	66.83	70.17	62.90	72.73	43.61	61.00	38.10	16.00	49.40	37.62
RoboPoint-13B [46]	75.85	77.83	44.5	61.29	60.84	44.36	69.9	21.90	41.30	49.31	20.79
<i>TIGeR</i>	93.85	96.33	95.17	91.94	<u>86.01</u>	60.15	82.11	82.86	32.79	80.82	70.30

also improve its geometry computation accuracy, powered by code generation and numerical calculation.

B. Simulation Evaluation for Robotics

We evaluate our model on the Open6DOR V2 Level 0 and Level 1 of position track, comparing against VLA-based baselines (pretrained Octo, LIBERO-finetuned OpenVLA) and SoFar, which integrates Florence-2, SAM, GPT-4o, and GSNet. This benchmark requires models to understand different spatial concepts (e.g., “left”, “front”, “between”) and manipulate objects. We find that, by predicting a single 3D target point compared to end-to-end VLA models or 2D bounding box-based methods, *TIGeR* mitigates 2D detection ambiguity under occlusion, yielding an 11.3% absolute improvement in success rate (Tab. II) against these baselines.

C. Real-world Evaluation for Robotics

Our real-world setting is shown in Fig. 6. We deploy the model on a UR5 robotic arm equipped with an Intel RealSense L515 RGB-D camera. We evaluate *TIGeR* on a suite of spatially-aware tasks that require geometry reasoning

TABLE II: Simulation results on Open6DOR V2 position track (Level 0 and Level 1). We report the success rate (%).

Part	Octo	OpenVLA	SoFar	<i>TIGeR</i>
Level 0	51.2	51.6	75.3	84.8
Level 1	12.7	13.1	65.6	81.0
Avg.	39.8	40.2	72.4	83.7

and high-precision 3D localization, which are infeasible for standard 2D pixel-based VLMs due to their inability to infer depth, camera extrinsics, or occluded/hovering positions. Specifically, we use AnyGrasp [47] to predict a valid grasp pose and further estimate the target’s 3D location in the camera frame to guide the placement.

We evaluate our method on four challenging real-world, spatially-aware manipulation tasks, with results summarized in Tab. III. These tasks require fine-grained spatial reasoning, such as metric-precision placement (e.g., “0.1m to the right”) and handling partial occlusions (e.g., the “back” of an object being invisible to the camera). Our method, *TIGeR*, consistently outperforms all baselines. It achieves a 55% success rate in metric-precision placement, demonstrating its ability

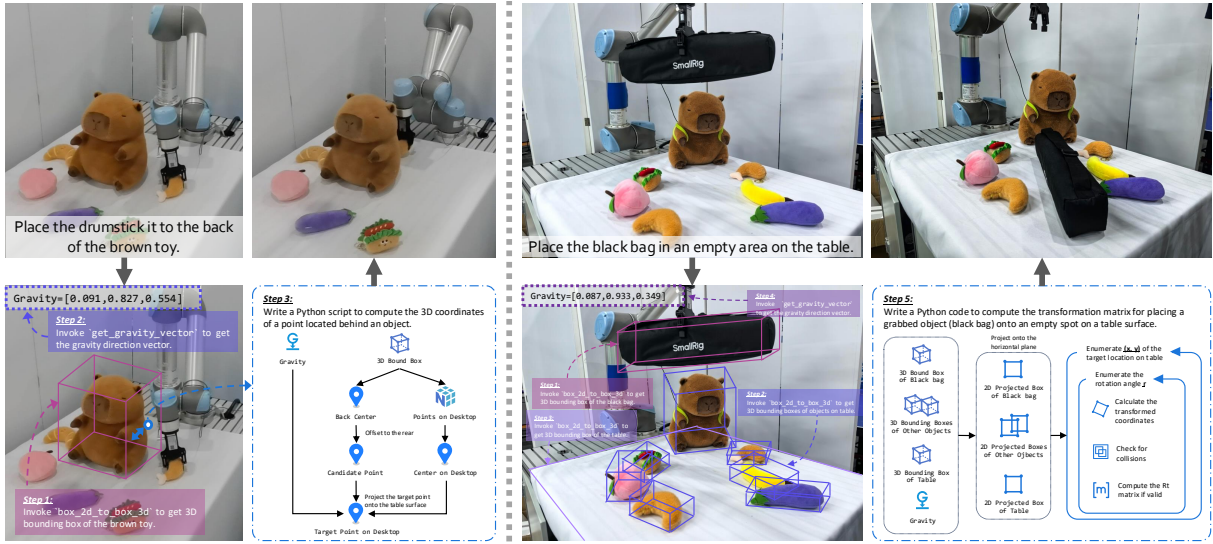


Fig. 5: Illustration of *TIGeR*'s tool-integrated reasoning in real-robot tasks.. “Pick up the drumstick and place it to the back of the brown toy” requires precise 3D spatial reasoning to localize the target position despite occlusion of the rear region. Similarly, “Place the black bag in an empty area on the table” demands complex geometric computation to identify feasible free space in cluttered scenes. We visualize the robot’s execution through sequential video frames, along with *TIGeR*'s step-by-step tool calls and their outputs, overlaid on the corresponding images.




Fig. 6: Real-world settings, including the UR5 robotic arm, L515 RGB-D camera, Inspire EG2-4B2 gripper. We show the captured RGB image with the depth map. TABLE III: Real-world evaluation of manipulation tasks requiring precise geometry reasoning. We report success rates (%). For the task “Pick up the hotdog and place it 0.1 m to the right of the peach”, success is defined as the actual placement distance being between 0.05 m and 0.15 m.

Manipulation tasks of 3D spatial positioning	Success Rate(%) \uparrow OpenVLA RoboPoint <i>TIGeR</i>		
Pick up the croissant and place it to the left of the drumstick.	0.00	65.00	75.00
Pick up the hotdog and place it 0.1 m to the right of the peach.	0.00	10.00	55.00
Pick up the drumstick and place it to the back of the brown toy (<i>facing occlusion of rear region</i>).	0.00	0.00	60.00
Pick up the bottle and move it above the plant	0.00	0.00	70.00

to accurately map spatial constraints to 3D coordinates. For spatial relations such as “back” and “above”, *TIGeR* achieves 60–70% success rates, significantly higher than baselines lacking depth-aware reasoning. Moreover, as shown in Fig. 5, *TIGeR* is the only method capable of resolving partial occlusions through precise geometric computation, highlighting the effectiveness of our explicit 3D coordinate representation

TABLE IV: Ablation studies on data recipes and reward design. We evaluate the impact of using template-based data (T.B.) vs. LLM-rewritten data (Rewritten.), and the individual effects of five reward functions. Tool. and Param. denote the tool invocation reward and parameter content reward. The last column reports the average accuracy (%) across geometric reasoning benchmarks described in Sec. IV-A.

Data Recipe		Reward					Avg.
T.B.	Rewritten.	Format	Tool.	Param.	Code	Answer	Acc.(%)
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	28.92
\checkmark	\times	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	74.46
\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	\checkmark	78.81
\checkmark	\checkmark	\checkmark	\times	\checkmark	\checkmark	\checkmark	76.12
\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark	\checkmark	75.96
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\times	\checkmark	79.11
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\times	78.39
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	79.30

for real-world manipulation in complex environments.

D. Ablation Study

We conduct comprehensive ablation studies to analyze the contribution of each component in our framework, as shown in Tab. IV. For the data recipe, we compare template-based synthetic data with LLM-rewritten data, finding that combining both achieves the best performance (79.30%), while using only rewritten data significantly degrades accuracy to 28.92%. This validates the importance of structured template data for establishing robust geometric reasoning foundations. For the reward design in RFT, we evaluate five reward components and find that removing any single reward leads to performance degradation. These results demonstrate that both diverse training data and comprehensive reward signals are essential for enabling reliable tool-integrated geometric reasoning in VLMs.

V. CONCLUSION

In this work, we present *TIGeR*, a novel tool-integrated framework that achieves precise geometric reasoning beyond spatial reasoning. In detail, we propose a two-stage SFT-RFT training pipeline with our proposed hierarchical reward design, enabling exact geometric computation via code generation and execution on calibrated metric inputs during the reasoning process. We also introduce *TIGeR-300K*, a large-scale dataset explicitly designed for geometric reasoning via programmatic tool calls, with a complete tool invocation sequence and intermediate computations. Extensive experiments demonstrate the superiority of the proposed approach and highlight its potential for a broad range of robotic applications requiring precise numerical computation.

This work was supported by the National Natural Science Foundation of China (62132001, 62476011), the Beijing Natural Science Foundation (L252060), the Fundamental Research Funds for the Central Universities, and the Xiaomi Young Talents Program/Xiaomi Foundation.

REFERENCES

- [1] Z. Fan *et al.*, “Vlm-3r: Vision-language models augmented with instruction-aligned 3d reconstruction,” *arXiv preprint arXiv:2505.20279*, 2025.
- [2] D. Wu *et al.*, “Spatial-mlm: Boosting mllm capabilities in visual-based spatial intelligence,” *arXiv preprint arXiv:2505.23747*, 2025.
- [3] D. Zheng *et al.*, “Learning from videos for 3d world: Enhancing mllms with 3d vision geometry priors,” *arXiv preprint arXiv:2505.24625*, 2025.
- [4] A.-C. Cheng *et al.*, “Spatialrgpt: Grounded spatial reasoning in vision-language models,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 135 062–135 093, 2024.
- [5] Y. Liu *et al.*, “Ssr: Enhancing depth perception in vision-language models via rationale-guided spatial reasoning,” *arXiv preprint arXiv:2505.12448*, 2025.
- [6] E. Zhou *et al.*, “Roborefer: Towards spatial referring with reasoning in vision-language models for robotics,” *arXiv preprint arXiv:2506.04308*, 2025.
- [7] W. Cai *et al.*, “Spatialbot: Precise spatial understanding with vision language models,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 9490–9498.
- [8] G. R. Team *et al.*, “Gemini robotics: Bringing ai into the physical world,” *arXiv preprint arXiv:2503.20020*, 2025.
- [9] P. Lu *et al.*, “Octotools: An agentic framework with extensible tools for complex reasoning,” *arXiv preprint arXiv:2502.11271*, 2025.
- [10] X. Li *et al.*, “Search-o1: Agentic search-enhanced large reasoning models,” *arXiv preprint arXiv:2501.05366*, 2025.
- [11] H. Tan *et al.*, “Robo-next: A unified memory-based framework for lifelong, scalable, and robust multi-robot collaboration,” *arXiv preprint arXiv:2510.26536*, 2025.
- [12] G. Dong *et al.*, “Self-play with execution feedback: Improving instruction-following capabilities of large language models,” *arXiv preprint arXiv:2406.13542*, 2024.
- [13] Z. Li *et al.*, “Do you have freestyle? expressive humanoid locomotion via audio control,” *arXiv preprint arXiv:2512.23650*, 2025.
- [14] B. R. Team *et al.*, “Robobrain 2.0 technical report,” *arXiv preprint arXiv:2507.02029*, 2025.
- [15] H. Tan *et al.*, “Robobrain 2.5: Depth in sight, time in mind,” *arXiv preprint arXiv:2601.14352*, 2026.
- [16] E. Zhou *et al.*, “Robotracer: Mastering spatial trace with reasoning in vision-language models for robotics,” *arXiv preprint arXiv:2512.13660*, 2025.
- [17] Z. Li *et al.*, “Robomirror: Understand before you imitate for video to humanoid locomotion,” *arXiv preprint arXiv:2512.23649*, 2025.
- [18] J. Feng *et al.*, “Retool: Reinforcement learning for strategic tool use in llms,” *arXiv preprint arXiv:2504.11536*, 2025.
- [19] X. Li *et al.*, “Torl: Scaling tool-integrated rl,” *arXiv preprint arXiv:2503.23383*, 2025.
- [20] Z. Li *et al.*, “From language to locomotion: Retargeting-free humanoid control via motion latent guidance,” *arXiv preprint arXiv:2510.14952*, 2025.
- [21] B. Chen *et al.*, “SpatialVlm: Endowing vision-language models with spatial reasoning capabilities,” in *CVPR*, 2024.
- [22] C. H. Song *et al.*, “Robospacial: Teaching spatial understanding to 2d and 3d vision-language models for robotics,” *CVPR*, 2025.
- [23] A. Ray *et al.*, “Sat: Spatial aptitude training for multimodal language models,” *arXiv*, 2024.
- [24] E. Zhou *et al.*, “Code-as-monitor: Constraint-aware visual programming for reactive and proactive robotic failure detection,” *arXiv preprint arXiv:2412.04455*, 2024.
- [25] Y. Wang *et al.*, “ π^3 : Scalable permutation-equivariant visual geometry learning,” *arXiv preprint arXiv:2507.13347*, 2025.
- [26] R. Wang and others, “Moge-2: Accurate monocular geometry with metric scale and sharp details,” *arXiv preprint arXiv:2507.02546*, 2025.
- [27] J. Wang *et al.*, “Vggt: Visual geometry grounded transformer,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 5294–5306.
- [28] N. Ravi *et al.*, “Sam 2: Segment anything in images and videos,” *arXiv preprint arXiv:2408.00714*, 2024.
- [29] A. Yang *et al.*, “Qwen3 technical report,” *arXiv preprint arXiv:2505.09388*, 2025.
- [30] J. Lazarow *et al.*, “Cubify anything: Scaling indoor 3d object detection,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 22 225–22 233.
- [31] S. Liu *et al.*, “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” in *European conference on computer vision*. Springer, 2024, pp. 38–55.
- [32] Y. Zhang *et al.*, “Recognize anything: A strong image tagging model,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 1724–1732.
- [33] B. Xiao *et al.*, “Florence-2: Advancing a unified representation for a variety of vision tasks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 4818–4829.
- [34] W. Hong *et al.*, “Glm-4.1 v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning,” *arXiv e-prints*, pp. arXiv–2507, 2025.
- [35] Z. Shao *et al.*, “Deepseekmath: Pushing the limits of mathematical reasoning in open language models,” *arXiv preprint arXiv:2402.03300*, 2024.
- [36] P. Tong *et al.*, “Cambrian-1: A fully open, vision-centric exploration of multimodal llms,” vol. 37, pp. 87 310–87 356, 2024.
- [37] X. Fu *et al.*, “Blink: Multimodal large language models can see but not perceive,” in *European Conference on Computer Vision*. Springer, 2024, pp. 148–166.
- [38] C. H. Song *et al.*, “Robospacial: Teaching spatial understanding to 2d and 3d vision-language models for robotics,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 15 768–15 780.
- [39] M. Du *et al.*, “Embspatial-bench: Benchmarking spatial understanding for embodied tasks with large vision-language models,” *arXiv preprint arXiv:2406.05756*, 2024.
- [40] Y.-H. Liao *et al.*, “Reasoning paths with reference objects elicit quantitative spatial reasoning in large vision-language models,” *arXiv preprint arXiv:2409.09788*, 2024.
- [41] J. Achiam *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [42] G. Team *et al.*, “Gemini: a family of highly capable multimodal models,” *arXiv preprint arXiv:2312.11805*, 2023.
- [43] Z. Liu *et al.*, “Nvila: Efficient frontier visual language models,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 4122–4134.
- [44] S. Bai *et al.*, “Qwen2. 5-vl technical report,” *arXiv preprint arXiv:2502.13923*, 2025.
- [45] M. Foutter *et al.*, “Space-llava: a vision-language model adapted to extraterrestrial applications,” *arXiv preprint arXiv:2408.05924*, 2024.
- [46] W. Yuan *et al.*, “Robopoint: A vision-language model for spatial affordance prediction for robotics,” *arXiv preprint arXiv:2406.10721*, 2024.
- [47] H.-S. Fang *et al.*, “Anygrasp: Robust and efficient grasp perception in spatial and temporal domains,” *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3929–3945, 2023.