One Agent to Guide Them All: Empowering MLLMs for Vision-and-Language Navigation via Explicit World Representation

Zerui Li[1,†] **Hongpei Zheng**[2,†] **Fangguo Zhao**[3] **Yanyan Shao**[4]
Jian Zhou[1] Sihao Lin[1,6] Shijie Li[5] Qi Wu[1,6,*]

[1]Australian Institute for Machine Learning, Adelaide University
[2]The University of Manchester [3]Zhejiang University
[4]Zhejiang Wanli University
[5]Agency for Science, Technology and Research (A*STAR)
[6]Responsible AI Research Centre,
Australian Institute for Machine Learning, Adelaide University

[†]Joint first authors [*]Corresponding author

Abstract: A navigation agent must understand high-level semantic instructions and reason about precise spatial structure. Multimodal Large Language Models (MLLMs), with their strong generalization, are a promising foundation for such agents. We propose GTA, a decoupled design that separates low-level spatial state estimation from high-level semantic planning. We introduce an explicit metric world representation that combines accumulated geometry with navigation history and exposes this information to the MLLM through a multimodal interface for decision-making. We further introduce counterfactual branch reasoning that checks anticipated outcomes against subsequent observations and revises navigation decisions when they disagree. The world representation grounds waypoint selection in reconstructed metric geometry. GTA achieves 48.8% Success Rate (SR) on the full R2R-CE val-unseen split and 46.2% on a 260-episode English RxR-CE val-unseen subset. Additional comparisons on shared R2R-CE episodes and mechanism ablations support the framework's effectiveness. We also demonstrate transfer to wheeled, quadrupedal, and aerial robots. Project page: https://zeruili22.github.io/GTA/

1 Introduction

Vision-and-Language Navigation (VLN) [1] requires an embodied agent to follow natural-language instructions in a 3D environment. Multimodal Large Language Models (MLLMs) [2, 3], with their strong semantic reasoning and open-world knowledge, are promising generalist VLN agents [4, 5]. Beyond costly fine-tuning that erodes pretrained knowledge, MLLM-centric *zero-shot* systems are gaining traction. These systems [4, 5, 6, 7] typically adopt a tightly coupled design, jointly performing spatial and semantic reasoning over long observation streams. However, such designs struggle in challenging scenarios along two dimensions. At the representation level, relying on egocentric RGB alone forces the MLLM to infer global spatial structure implicitly, often hallucinating layouts (e.g., nonexistent room connections). These errors then propagate into semantic reasoning, yielding instruction-misaligned decisions. At the reasoning level, processing everything in a single step leaves little room for the flexible, multi-step strategies needed to fully elicit MLLM capabilities.

To address both, we propose GTA (Guide Them All), a zero-shot VLN framework that decouples low-level spatial state estimation from high-level semantic reasoning (Fig. 1). An explicit metric world representation, kept separate from the reasoning policy, grounds decisions in a consistent global layout so spatial errors no longer propagate. This separation in turn lets each part follow its own reasoning strategy, which a tightly coupled system cannot. Crucially, these are architectural components. We pair an external, queryable world representation with a stateful reasoning procedure

10th Conference on Robot Learning (CoRL 2026), Austin TX, USA.



Figure 1: **Bridging high-level reasoning and embodied execution with GTA.** *Left*: Existing MLLM-based VLN agents reduce the rich 3D environment into an oversimplified linearized text memory according to human knowledge. *Right*: Our GTA framework decouples spatial modeling from semantic reasoning and maintains rich spatial and historical information, enabling the MLLM to interact with it for decision-making.

over it, rather than reformatting observations into text, and validate their effect by ablations rather than prompt tuning. Unlike prior work relying on passive memory buffers or heuristic maps [8, 9, 10, 11], our world representation lets the agent actively query task-relevant information on demand, grounding decisions in an explicit layout and mitigating the hallucinations that plague egocentric reasoning. Building on this representation, we introduce a counterfactual reasoning strategy that further elicits MLLM capability. Rather than relying on a fixed waypoint predictor [12], the MLLM evaluates each candidate action by querying the world representation, reasoning about what each action would entail against a decomposed goal list, so the chosen action is both globally consistent and physically valid (Fig. 1).

We evaluate GTA on multiple VLN benchmarks [1, 13, 14], where it achieves state-of-the-art zero-shot performance across datasets and MLLM backbones [3, 15, 16, 17], and transfers to real robots. Our contributions are as follows:

- We propose GTA, a decoupled zero-shot VLN framework that mitigates error propagation and enables task-specific reasoning, achieving state-of-the-art zero-shot performance.
- We introduce an explicit world representation that gives the MLLM global metric context for decision-making, avoiding human-designed priors and passive memory summaries.
- We develop a counterfactual reasoning strategy operating directly on the world representation, enabling physically grounded reasoning that further elicits MLLM capability.

2 Related Work

Vision-and-Language Navigation (VLN): VLN requires an agent to follow natural-language instructions through a visual environment. Early benchmarks were mostly discrete [1, 18, 14], reducing navigation to selecting nodes on a predefined graph. Recent work shifts to continuous environments (VLN-CE) [13] for closer alignment with real robotics: replacing node “teleporting” with low-level actions sharply degrades traditional models [13, 19, 20, 21], as the agent must now reason about fine-grained geometry, track its history, and act in a continuous space. Bridging high-level linguistic reasoning and low-level continuous control thus remains a core bottleneck.

From Expert Models to Embodied Large Models: Early work tackled continuous control with task-specific expert models, notably waypoint-based architectures [22, 23, 24, 11, 25, 26]. These improved success on benchmarks like R2R-CE but act as “in-domain” experts tightly coupled to their training distribution, generalizing poorly across scenes. The rise of MLLMs shifted the paradigm toward embodied generalists: representative works adapt LLMs to navigation by mapping visual-linguistic features into the embedding space and adding navigation-specific tokens, combining in-domain proficiency with open-vocabulary reasoning [27, 28, 29]. Yet this adaptation is resource-intensive and, more critically, structurally couples the agent to a fixed backbone, leaving it unable to track rapidly improving foundation models [29].

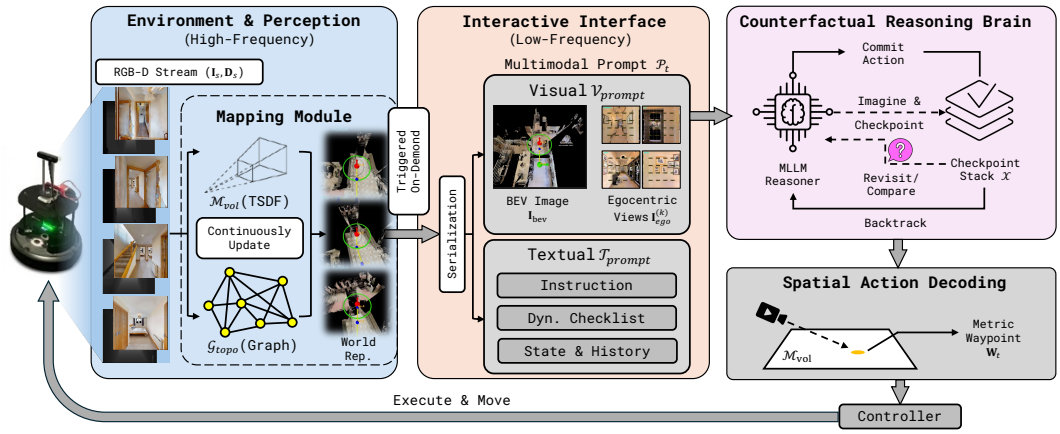


Figure 2: **Overview of the GTA Framework.** The *Mapping Module* (left) fuses RGB-D streams into a TSDF volume and topological graph, forming our world representation. The *Interactive Interface* (middle) serializes it on demand into a multimodal prompt with BEV/egocentric views, instruction, dynamic checklist, state, and history. The *Counterfactual Reasoning Brain* (right), driven by a frozen MLLM, forms branch-specific expectations, checks them against subsequent observations, and selects actions decoded into metric waypoints.

Zero-Shot Embodied Agents in VLN: These limitations motivated zero-shot agents that exploit the emergent reasoning of foundation models without task-specific training. Early efforts showed LLMs can reason topologically [8, 30] but stayed in discrete settings. In VLN-CE, contemporary zero-shot planners represent the environment by translating observations and topology into serialized text or language-based mental maps [6, 31, 32, 33]. Such language-centric representations preserve a high-level narrative of progress but lack an explicit, geometrically grounded world model [7, 12, 34], so under limited line-of-sight or cluttered geometry the agents struggle to resolve spatial ambiguity or predict precise waypoints. This gap motivates our explicit metric world representation. A complementary line studies 3D scene graphs: Hydra [35] constructs hierarchical spatial representations online, SayPlan [36] combines scene-graph search with path planning and feedback-based replanning, and structured query interfaces [37] retrieve task-relevant graph information for language grounding. GTA connects online metric geometry to continuous instruction following through multimodal prompting, pixel-to-waypoint grounding, and observation-based verification of branch expectations.

3 Methodology

We propose GTA, a hierarchical navigation framework that decouples low-level spatial state estimation from high-level semantic planning (Fig. 2). Instead of forcing one model to track geometry and reason about language at once, GTA maintains an explicit metric world representation $\mathcal{W}_t$ and lets a frozen MLLM reason over it. At each step $\mathcal{W}_t$ is serialized into a multimodal prompt, the MLLM picks a target, and the target is decoded into a metric waypoint. Perception runs continuously while reasoning fires on demand, so decisions rest on an accumulated map rather than a single noisy frame. We detail the mapping, interface, and reasoning modules in Algorithm 1, App. B.

3.1 Problem Setup

We address Vision-and-Language Navigation in Continuous Environments (VLN-CE). An agent placed in a continuous 3D environment receives a natural-language instruction $\mathcal{I}$, and must reach a target region within a success threshold from its start pose. Its pose at step t is $\mathbf{s}_t = (x_t, y_t, \theta_t)$ on the navigable floor plan, with planar position $\mathbf{p}_t = (x_t, y_t)$ and heading θ_t. The environment is only partially observed through an asynchronous RGB-D stream. At each reasoning step the policy emits a structured action $a_t = (m_t, u_t, v_t)$, comprising a view index m_t and a target pixel (u_t, v_t) within that view, later decoded into a metric waypoint (Sec. 3.5). We assume neither a pre-built map nor

a reward signal. The policy is realized by a frozen, off-the-shelf MLLM $\mathcal{F}$ reasoning on our world representation $\mathcal{W}_t$.

3.2 Explicit Metric World Representation

To make the agent aware of both spatial and historical context, the Metric Mapping Module asynchronously builds a hybrid representation $\mathcal{W}_t = \langle \mathcal{M}_{vol}, \mathcal{G}_{topo} \rangle$, comprising a volumetric map $\mathcal{M}_{vol}$ that encodes geometry from high-frequency RGB-D frames and a topological graph $\mathcal{G}_{topo}$ that encodes the visited trajectory. $\mathcal{W}_t$ accumulates spatial and trajectory context throughout navigation and provides the spatial input to the interface (Sec. 3.3).

Volumetric mapping ($\mathcal{M}_{vol}$). To accumulate dense metric geometry across frames, we fuse depth into a TSDF [38, 39]. A pixel $\mathbf{u}$ at depth d in frame $\mathcal{O}_s$ (yaw θ_s) is back-projected by $\pi^{-1}(\mathbf{u}, d) = d\,\mathbf{K}^{-1}[u, v, 1]^\top$ and mapped to the world frame via the body pose $\mathbf{T}_{wb}(s) \in SE(3)$ and extrinsics $\mathbf{T}_{bc}(\theta_s)$, giving world point $\mathbf{p}_w$ and camera center $\mathbf{p}_{cam}(s)$:

$$\mathbf{p}_w = \mathbf{T}_{wb}(s)\,\mathbf{T}_{bc}(\theta_s)\,\big[\pi^{-1}(\mathbf{u}, d);\,1\big], \qquad \mathbf{p}_{cam}(s) = \mathbf{T}_{wb}(s)\,\mathbf{T}_{bc}(\theta_s)\,[0,0,0,1]^\top. \tag{1}$$

For a voxel $\mathbf{v}$ at $\mathbf{x}$, the truncated projective signed distance is $\psi_t(\mathbf{x}) = \mathrm{clip}\big(d - \|\mathbf{x} - \mathbf{p}_{cam}(s)\|, -\mu, \mu\big)$ with truncation radius μ, fused with a confidence weight $\omega_t \in [0, 1]$ (decaying with depth and grazing angle) by a running average that accumulates into $\Omega_t(\mathbf{v})$:

$$S_t(\mathbf{v}) = \frac{\Omega_{t-1}(\mathbf{v})\,S_{t-1}(\mathbf{v}) + \omega_t\,\psi_t(\mathbf{x})}{\Omega_{t-1}(\mathbf{v}) + \omega_t}, \qquad \Omega_t(\mathbf{v}) = \Omega_{t-1}(\mathbf{v}) + \omega_t. \tag{2}$$

Its zero-level set defines the reconstructed scene surface, which is reused for BEV rendering (Sec. 3.3) and waypoint ray-casting (Eq. (9)).

Topological abstraction ($\mathcal{G}_{topo}$). The TSDF captures geometry but not where the agent has been. For long-horizon memory we add a topological graph $\mathcal{G}_{topo} = (\mathcal{V}, \mathcal{E}_{topo})$, where each node $v_i = \langle \mathbf{p}_i, c_i \rangle$ stores a metric position $\mathbf{p}_i \in \mathbb{R}^2$ and a visit count c_i, and $\mathcal{E}_{topo}$ links consecutive nodes. On reaching $\mathbf{p}_t$, ClusterNode merges the visit into the nearest existing node if within δ_{merge}, or creates a new node otherwise:

$$\mathcal{V} \leftarrow \begin{cases} \mathrm{Update}(v_{nearest}) & d_{min} < \delta_{merge} \\ \mathcal{V} \cup \{\mathrm{Node}(\mathbf{p}_t)\} & \text{otherwise,} \end{cases} \tag{3}$$

where $d_{min} = \min_{v_i \in \mathcal{V}} \|\mathbf{p}_t - \mathbf{p}_i\|_2$ and $\delta_{merge} = 0.8\,\mathrm{m}$. Here $\mathrm{Update}(v_{nearest})$ denotes the node set with only the nearest node's visit count incremented; its position remains unchanged. $\mathrm{Node}(\mathbf{p}_t)$ creates a node at $\mathbf{p}_t$ with initial visit count one. The visit counts c_i make loops and backtracking explicitly detectable, feeding the safety prompt and counterfactual revision.

3.3 Interactive Reasoning Interface

An MLLM cannot consume $\mathcal{W}_t$ directly, so a training-free interface ϕ serializes it into a multimodal prompt $\mathcal{P}_t = \{\mathcal{V}_{prompt}, \mathcal{T}_{prompt}\}$ with visual and textual parts. The interface is non-parametric; Sec. 3.6 later replaces it with a learnable ϕ_θ.

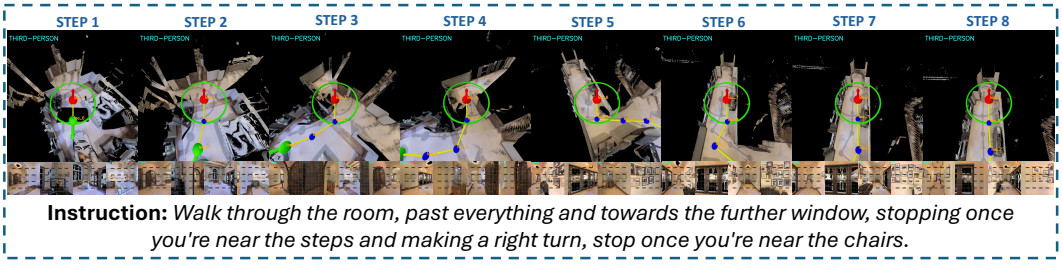
Egocentric view selection. To get 360° context without extra actuation, we select four egocentric views from past frames rather than capturing new ones. From the pool $\mathcal{N}_t$ of the K=20 spatially nearest past frames, for each target angle $\alpha_k \in \{0°, 90°, 180°, 270°\}$ we pick the frame whose orientation best matches $\mathbf{q}_t \otimes \mathbf{q}(\alpha_k)$ by absolute quaternion similarity:

$$s_k = \arg\max_{s \in \mathcal{N}_t} \big|\langle\, \mathbf{q}_s,\ \mathbf{q}_t \otimes \mathbf{q}(\alpha_k)\,\rangle\big|, \qquad \big(\mathbf{I}^{(k)}_{ego}, \mathbf{D}^{(k)}_{ego}\big) = \big(\mathbf{I}_{s_k}, \mathbf{D}_{s_k}\big), \tag{4}$$

each frame assigned to at most one direction, empty if none.

Visual prompting with a coordinate grid. So the MLLM can emit precise targets without a regression head, we render the TSDF surface into a bird's-eye view $\mathbf{I}_{bev}$ and overlay an integer coordinate grid $[0, 1000]^2$ on it and the egocentric views, giving

$$\mathcal{V}_{prompt} = \big[\mathbf{I}_{bev},\ \mathbf{I}^{(1)}_{ego}, \ldots, \mathbf{I}^{(4)}_{ego}\big]. \tag{5}$$



Instruction: *Walk through the room, past everything and towards the further window, stopping once you're near the steps and making a right turn, stop once you're near the chairs.*

Figure 3: **Qualitative visualization of a navigation episode in R2R-CE.** The top row displays the agent's world representation (top-down metric view), while the bottom row shows the four selected egocentric views at each step. The planned trajectory is marked in yellow, with blue dots indicating waypoints and the red arrow showing the agent's current pose.

The two modalities are complementary. The egocentric RGB views carry appearance for open-vocabulary grounding of entities named in $\mathcal{I}$, while the BEV grid carries metric structure for precise waypoint coordinates.

Procedural reasoning blueprint ($\mathcal{T}_{prompt}$). The textual half is assembled by concatenation ($\oplus$):

$$\mathcal{T}_{prompt} = \mathcal{C}_t \oplus \text{State}(\mathcal{G}_{topo}, \mathcal{S}_{alert}) \oplus \text{History}(\mathcal{H}_t) \oplus \mathcal{I}, \tag{6}$$

with the following fields. **Dynamic checklist $\mathcal{C}_t$:** the instruction is decomposed into a checklist (e.g. "[x] reach the door", "[ ] turn left") that the MLLM updates each step so the immediate action stays aligned with the high-level plan. **Topological & physical state** $\text{State}(\mathcal{G}_{topo}, \mathcal{S}_{alert})$**:** the agent's situation in the graph plus two safety signals collected in $\mathcal{S}_{alert}$, namely vertical awareness ("upstairs"/"downstairs" when $|\Delta h| > 0.3\,\text{m}$) and alerts raised when the previous action failed or when a loop is detected (a node visited at least c_{loop}=3 times within a six-step window, or oscillation between two nodes). **Execution history** $\text{History}(\mathcal{H}_t)$**:** $\mathcal{H}_t$ is the ordered log of recent (thought, view, action) tuples, truncated to a sliding window of length H=20. **Global instruction $\mathcal{I}$:** the original command, kept immutable. The assembled $\mathcal{P}_t$ is the sole input to the reasoner in Sec. 3.4.

3.4 Counterfactual Reasoning Brain

Given $\mathcal{P}_t$, the frozen MLLM $\mathcal{F}$ acts as the high-level policy. A greedy reasoner that commits to one direction per step cannot recover from an early wrong branch. We therefore let the agent imagine the outcome of each branch before committing and revisit it when later evidence disagrees.

Specifically, $\mathcal{F}$ uses the egocentric images and rendered map to identify directions that appear traversable and relevant to the current checklist item $\mathcal{C}_t$. We denote the resulting viable set by $\mathcal{B}_t$ and treat the step as a branch point when $|\mathcal{B}_t| > 1$. At such a point, $\mathcal{F}$ emits a per-direction expectation $e_t^{(k)}$ (e.g., "a hallway to a kitchen on the right") for each $k \in \mathcal{B}_t$ and pushes a checkpoint onto a stack $\mathcal{K}$,

$$\kappa_t = \big\langle\, \mathbf{p}_t,\ \{(k, e_t^{(k)})\}_{k\in\mathcal{B}_t},\ \mathcal{B}_t^{\text{blk}} \,\big\rangle, \qquad \mathcal{B}_t^{\text{blk}} = \emptyset, \tag{7}$$

where $\mathcal{B}_t^{\text{blk}}$ records directions already ruled out. The initially selected branch is $k^\star = m_t$, the view index of the action emitted at that checkpoint. After advancing along $k^\star$, $\mathcal{F}$ judges, in natural language, whether the observation matches the stored expectation $e_t^{(k^\star)}$. A rising visit count c_i or a loop alert from $\mathcal{G}_{topo}$ (Sec. 3.3) serves as additional evidence that the branch is unproductive.

The reasoning becomes counterfactual when the branch is judged inconsistent. The agent then retrieves the latest active checkpoint $\kappa_{t'}$ and physically returns to $\mathbf{p}_{t'}$ by retracing visited nodes in $\mathcal{G}_{topo}$ with the same local planner used for forward motion. It blocks the failed branch and re-decides over the untried directions,

$$\mathcal{B}_{t'}^{\text{blk}} \leftarrow \mathcal{B}_{t'}^{\text{blk}} \cup \{k^\star\}, \qquad k_{t'} = \arg\max_{k\in\mathcal{B}_{t'}\setminus\mathcal{B}_{t'}^{\text{blk}}} g\big(k \mid \mathcal{P}_{t'}, \mathcal{I}\big), \tag{8}$$

where $g(\cdot)$ is $\mathcal{F}$'s preference over the remaining directions. Explicit blocking prevents re-selecting a rejected branch at that checkpoint. Once a checkpoint's directions are exhausted ($\mathcal{B}_{t'} \setminus \mathcal{B}_{t'}^{\text{blk}} = \emptyset$),

it is released and a safety alert is injected into $\mathcal{T}_{prompt}$ to backtrack one level further up $\mathcal{K}$ or replan; a branch confirmed consistent retires its checkpoint. This limits repeated exploration at each checkpoint while allowing recovery from an early wrong branch.

3.5 Spatial Action Decoding and Execution

The MLLM returns an action $a_t = (m_t, u_t, v_t)$, which we turn into a metric goal by ray-casting from the camera $\mathbf{p}_{cam}^{(m_t)}$ along the viewing ray $\mathbf{r}^{(m_t)}(u_t, v_t) = \mathbf{K}^{-1}[u_t, v_t, 1]^\top$ (direction only, distinct from π^{-1}), transformed to the world frame by the view's rotation before being cast against the TSDF surface:

$$\mathbf{w}_t = \mathrm{RayCast}\big(\mathbf{p}_{cam}^{(m_t)},\, \mathbf{r}^{(m_t)}(u_t, v_t),\, \mathcal{M}_{vol}\big), \tag{9}$$

thereby grounding the semantic decision in reconstructed metric geometry. A deterministic local planner attempts to reach $\mathbf{w}_t$; execution failures and loop alerts are passed to the reasoner on the next step through the safety mechanism (Sec. 3.3).

3.6 GTA-Light: Extending to a Learnable World Interpreter

Lightweight neural interpreter. The zero-shot system relies on a strong frozen MLLM for high-level reasoning. Proprietary API-based models (e.g., GPT-5.1 [16]) require remote access, while large open-source models (e.g., Qwen3-VL-235B [3]) impose substantial compute requirements. To reduce these deployment costs, we learn a lightweight interpreter ϕ_θ, instantiated on a compact open-source MLLM (InternVL3-8B [15]), that maps the same world state to the same prompt space, $\hat{\mathcal{P}}_t = \phi_\theta(\mathcal{W}_t, \mathcal{I})$. We refer to the full system driven by this distilled interpreter as GTA-Light. The reasoning protocol is unchanged; only the translator becomes trainable.

Let $\mathcal{P}_t^\star = (y_1, \ldots, y_L)$ be the target prompt from a teacher MLLM $\mathcal{F}_{tea}$. We fine-tune ϕ_θ with a supervised sequence-modeling loss over the target tokens,

$$\mathcal{L}_{\text{sft}}(\theta) = -\sum_{(\mathcal{W}_t, \mathcal{P}_t^\star)\in\mathcal{D}} \sum_{\ell=1}^{L} \log p_\theta\big(y_\ell \mid y_{<\ell}, \mathcal{W}_t, \mathcal{I}\big), \tag{10}$$

and a task-aware decoding loss on the structured action fields (view index m_t, target pixel (u_t, v_t)) to sharpen action accuracy,

$$\mathcal{L}_{\text{act}} = \lambda_{\text{view}}\,\mathrm{CE}(\hat{m}_t, m_t^\star) + \lambda_{\text{coord}}\big\|(\hat{u}_t, \hat{v}_t) - (u_t^\star, v_t^\star)\big\|_1, \tag{11}$$

giving $\min_\theta\ \mathcal{L}_{\text{sft}} + \mathcal{L}_{\text{act}}$. For efficiency we use LoRA [40], freezing $\mathbf{W}_0$ and learning a low-rank residual $\mathbf{W} = \mathbf{W}_0 + \mathbf{BA}$ with $\mathrm{rank}(\mathbf{BA}) \ll \min(d_{in}, d_{out})$. QLoRA [41] further back-propagates through a frozen 4-bit backbone. Training data are generated from the R2R training split by replaying A* trajectories through the world representation and querying the teacher. No validation scene, episode, or instruction is used for data construction or student training. We apply input- and prompt-level robustness augmentations $\mathcal{A}_\epsilon^W, \mathcal{A}_\eta^P$ so that the objective becomes

$$\min_\theta\ \mathbb{E}_{(\mathcal{W}_t, \mathcal{P}_t^\star)\sim\mathcal{D}}\, \mathbb{E}_{\epsilon,\eta}\Big[\mathcal{L}\big(\phi_\theta\big(\mathcal{A}_\epsilon^W(\mathcal{W}_t),\, \mathcal{I}\big),\, \mathcal{A}_\eta^P(\mathcal{P}_t^\star)\big)\Big]. \tag{12}$$

Full data construction and augmentation details are in Appendix C.

4 Experiments

4.1 Experimental Setup

We evaluate on two continuous-environment benchmarks, R2R-CE [13] and RxR-CE [47], in the Habitat simulator [48], using the standard VLN-CE metrics (SR, OSR, SPL, NE, nDTW and success within 3 m). Unless specified, the reasoner is GPT-5.1. We further study backbone scalability in Sec. 4.3. The main results use the full R2R-CE val-unseen split (1,839 episodes) and a 260-episode English RxR-CE val-unseen subset. The GTA results in Table 1 are from a representative run near the median of three runs, with observed spreads of ± 0.8 percentage points in SR and ± 1.1 in SPL. The backbone and procedural-blueprint ablations use a curated 180-episode R2R-CE subset. The

Table 1: **Comparison in continuous environments on R2R-CE and RxR-CE Val-Unseen splits.** The **best supervised** results are in **bold**, and the best zero-shot results are underlined. We report the full validation split for R2R-CE (1839 episodes) and a sampled subset for RxR-CE (260 episodes).

#	Methods	R2R-CE					RxR-CE			
		NE↓	OSR↑	SR↑	SPL↑	nDTW↑	NE↓	SR↑	SPL↑	nDTW↑
Supervised Learning:										
1	CMA [42]	6.30	49.0	38.0	33.0	–	10.4	24.1	19.1	37.4
2	VLN-BERT [42]	5.74	53.0	44.0	39.0	–	8.98	27.1	22.7	46.7
3	GridMM [11]	5.11	61.0	49.0	41.0	–	–	–	–	–
4	ETPNav [25]	4.71	65.0	57.0	49.0	–	5.64	54.8	44.9	61.9
5	BEVBert [26]	4.70	67.0	59.0	50.0	–	4.80	64.4	–	65.4
6	HNR [43]	4.42	67.0	61.0	51.0	–	5.51	56.4	46.7	63.6
7	NavFoM [29]	4.61	72.1	61.7	55.3	–	4.74	64.4	**56.2**	65.8
8	Efficient-VLN [44]	**4.18**	**73.7**	**64.2**	**55.9**	–	**3.88**	**67.0**	54.3	**68.4**
Zero-Shot Learning:										
9	Open-Nav [7]	6.70	23.0	19.0	16.1	45.8	–	–	–	–
10	CA-Nav [33]	7.58	48.0	25.3	10.8	–	10.37	19.0	6.0	–
11	Smartway [12]	7.01	51.0	29.0	22.5	–	–	–	–	–
12	STRIDER [45]	6.91	39.0	35.0	30.3	51.8	11.19	21.2	9.6	30.1
13	VLN-Zero‡ [46]	5.97	51.6	42.4	26.3	–	9.13	30.8	19.0	–
14	BZS-VLN [34]	6.12	55.0	41.0	25.4	–	7.56	35.7	21.7	42.4
15	**GTA-Light (Ours)**†	5.41	53.9	45.1	39.6	58.8	7.33	42.5	36.2	54.1
16	**GTA (Ours)**	4.95	56.2	48.8	41.8	60.4	6.29	46.2	39.3	57.4

† GTA-Light is the lightweight variant (Sec. 3.6), distilled into InternVL3-8B, it is trained and thus excluded from the zero-shot comparison.
‡ VLN-Zero pre-explores each scene to construct its map before evaluation; its results are included for reference under a different map-access setting.

additional R2R-CE comparisons and mechanism ablations in Sec. 4.4 use Open-Nav's released 100-episode set [7, 12, 45, 49]. We compare against supervised (#1–#8) and zero-shot (#9–#14) baselines in Tab. 1. Full benchmark, metric, baseline, and subset details are in Appendix A.

4.2 Main Results

We present a quantitative comparison of our proposed GTA framework against state-of-the-art baselines in R2R-CE and RxR-CE, as summarized in Tab. 1. The results show that world representation substantially improves the navigation performance of generalist MLLMs.

State-of-the-Art Zero-Shot Performance As evidenced in Tab. 1, GTA (#16) sets a new zero-shot state-of-the-art on both R2R-CE and RxR-CE, surpassing the prior best zero-shot agent BZS-VLN [34] (#14) by 7.8% SR and 16.4% SPL on R2R-CE. These results use the same evaluation split but different backbones; Sec. 4.4 additionally reports shared-episode comparisons including STRIDER with GPT-5.1. The large SPL gain indicates that GTA navigates with map-informed efficiency rather than exploratory guessing. By decoupling semantic reasoning from spatial execution and grounding decisions in an explicit metric world representation, the MLLM produces precise, physically valid waypoints instead of coarse directional cues. Notably, these gains do not require a frontier backbone. GTA-Light, the same framework driven by a distilled 8B interpreter (Sec. 3.6), retains most of the performance (45.1% SR / 39.6% SPL on R2R-CE, vs. 48.8% / 41.8%) and still surpasses all prior zero-shot methods. This indicates that the improvements stem from the framework itself rather than from the scale of the backbone MLLM.

Competitive with Supervised Experts Although GTA operates in a strictly zero-shot manner without any parameter updates, it performs comparably to several supervised baselines. As shown in Tab. 1, GTA surpasses classic supervised agents such as CMA [13] (#1) and VLN-BERT [42] (#2), and narrows the gap to recent supervised methods like Efficient-VLN [44] (#8). These results suggest that much of the navigation competence needed for VLN already resides in off-the-shelf MLLMs, and that a metric-grounded framework is a promising, complementary direction for eliciting it, rather than a replacement for in-domain training.

Robustness on Complex Instructions (RxR-CE). RxR-CE features longer trajectories and fine-grained multilingual instructions. Prior zero-shot agents degrade sharply here. BZS-VLN [34] (#14) drops from 41.0% SR on R2R-CE to 35.7% on RxR-CE, whereas GTA (#16) sustains 46.2% SR (vs. 48.8% on R2R-CE), and even the distilled GTA-Light (#15) holds 42.5%, still above the prior

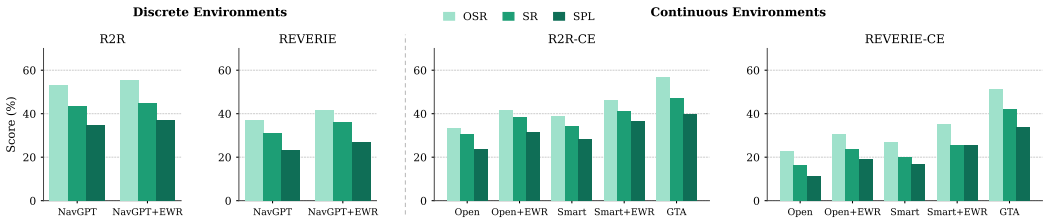


Figure 4: **Effect of the World Representation across environments.** Each panel reports OSR/SR/SPL. Adding WR (world representation) improves every baseline, and GTA leads on all metrics in continuous environments. Discrete: R2R / REVERIE; continuous: R2R-CE / REVERIE-CE. Within each continuous-environment comparison, Open-Nav [7], SmartWay [12], and GTA use the same evaluation episodes and the GPT-5.1 reasoner.

Table 2: **Ablation on MLLM backbones.**

Backbone	R2R-CE Val-Sampled (180)				
	NE↓	OSR↑	SR↑	SPL↑	nDTW↑
Qwen3-VL-235B	5.98	46.1	37.2	29.5	51.3
Gemini-2.5 Pro	5.35	48.7	42.3	34.4	**56.2**
GPT-5.1	**5.27**	**56.7**	**47.2**	**39.6**	55.8

Table 3: **Ablation on Procedural Reasoning Blueprints (PB).**

Method	R2R-CE Val-Sampled (180)				
	NE↓	OSR↑	SR↑	SPL↑	nDTW↑
GTA (w/o PB)	5.38	48.3	45.0	38.1	52.7
GTA (w/ PB)	**5.27**	**56.7**	**47.2**	**39.6**	**55.8**

best. This resilience indicates that the world representation provides a geometric anchor that buffers against instructional complexity. An additional 300-episode evaluation is reported in Sec. 4.4.

4.3 Ablation Study

Unless noted otherwise, ablations are conducted on the curated 180-episode R2R-CE subset (Sec. 4.1), so absolute numbers differ from the full-split results in Tab. 1.

Impact of World Representation. To verify the necessity of our metric mapping module, we equip baseline agents with and without world representation. As shown in Fig. 4, world representation yields consistent gains across both discrete and continuous environments, supporting our hypothesis (Sec. 1) that a structured, metric-accurate description of space helps MLLMs ground semantic decisions rather than rely on reactive visual matching. The effect is most pronounced on REVERIE and REVERIE-CE, whose coarse-grained object goals (e.g., "find the pillow in the bedroom") lack intermediate directional cues. Implicit agents degrade sharply on such tasks for want of a persistent spatial memory, whereas our agents remain robust by reasoning over the mapped structure to infer room connections and unexplored frontiers. This indicates that an explicit spatial representation (world representation) is particularly effective for handling complex, large-scale environments.

Scalability with MLLM Reasoning Capabilities. By decoupling spatial estimation from decision-making, GTA is model-agnostic and can directly adopt stronger backbones without architectural changes. As shown in Tab. 2, navigation performance scales with backbone reasoning ability, improving from 37.2% SR (Qwen3-VL-235B) to 42.3% (Gemini-2.5 Pro) and 47.2% (GPT-5.1). This suggests that GTA casts continuous navigation as a high-level reasoning problem, so it stands to benefit as MLLMs improve.

Impact of Procedural Reasoning Blueprints. The w/o PB variant removes exactly the dynamic checklist $\mathcal{C}_t$ and the State field from $\mathcal{T}_{prompt}$, while retaining the instruction, history, and visual inputs. As shown in Tab. 3, including PB increases SR from 45.0% to 47.2% and OSR from 48.3% to 56.7%, gains of 2.2 and 8.4 percentage points, respectively. The OSR increase indicates that the agent reaches the target vicinity more often, even when it does not stop successfully.

4.4 Additional Comparisons and Mechanism Ablations

Additional comparisons. On Open-Nav's released 100-episode R2R-CE set, GTA achieves 47.0% SR and 38.8% SPL (Tab. 4). Running STRIDER with GPT-5.1 on the same episodes yields 40.0% SR and 36.5% SPL, leaving gains of 7.0 and 2.3 percentage points. On an additional random sam-

Table 4: Additional comparisons (SR/SPL). R2R-CE uses the episode sets indicated below. RxR-CE uses a separate 300-episode sample without episode-level matching to BZS-VLN.

Protocol	Baseline		GTA
R2R-CE, 100-ep	STRIDER	35.0/30.3	47.0/38.8
w/ GPT-5.1	STRIDER	40.0/36.5	47.0/38.8
R2R-CE, 1,839-ep	BZS-VLN	41.0/25.4	48.8/41.8
RxR-CE, 300-ep	BZS-VLN	35.7/21.7	46.0/37.9

Table 5: Incremental counterfactual-reasoning ablations on the same 100-episode R2R-CE set; each row adds one mechanism, and (d) is the full GTA.

Variant	SR↑	SPL↑	nDTW↑
(a) Greedy, no backtracking	31.0	27.1	34.7
(b) + loop-alert backtrack	37.0	31.8	45.6
(c) + checkpoints, blocking	39.0	35.3	49.3
(d) + expectation checks	**47.0**	**38.8**	**53.6**

ple of 300 English RxR-CE episodes, GTA achieves 46.0% SR and 37.9% SPL. Since BZS-VLN's episode IDs were unavailable, its published results provide a reference comparison without controlling for episode composition.

Counterfactual reasoning. Table 5 separates three recovery mechanisms. Variant (a) acts greedily without backtracking. Variant (b) returns to the last junction on a loop alert but does not remember failed branches. Variant (c) adds checkpoints and blocked-direction sets, while detecting failure only through loop or collision signals. Full GTA (d) additionally checks branch-specific expectations against subsequent observations, allowing semantic inconsistencies to trigger rejection before a loop alert. The successive SR gains are 6.0, 2.0, and 8.0 percentage points. Component ablations are reported in Sec. F.

4.5 Real-World Deployment

To validate the practical utility of GTA, we deploy it on three physical platforms, a wheeled robot, a quadruped, and a quadrotor, in unseen real-world environments (hardware and protocol in Appendix D; qualitative rollouts in Appendix Fig. 5). As shown in Tab. 6, the visual domain gap between simulation and reality cripples end-to-end methods (VLN-BERT 16%, RDP 20%), whereas GTA sustains 40–42% SR across all three embodiments. We attribute this to our world representation. Although the MLLM still observes egocentric RGB, its decisions are anchored on the synthesized top-down metric map and topological graph, whose structure is domain-invariant (a wall is a wall in both simulation and reality), keeping the planner effective and stable across bodies despite the appearance shift. Reasoning latency remains compatible with real-time control under our asynchronous design (Appendix E).

Table 6: **Real-world navigation performance. Bold** = Ours.

Method	**SR↑**	**NE↓**
Supervised:		
VLN-BERT [42]	16.0	5.36
RDP [50]	20.0	5.45
Zero-Shot:		
SmartWay [12]	32.0	4.85
GTA – Wheeled	**40.0**	**3.66**
GTA – Quadruped	**40.0**	**3.59**
GTA – Drone	**42.0**	**3.50**

5 Conclusion and Limitation.

Conclusion. & Limitation. We present **GTA**, a decoupled zero-shot VLN framework that grounds MLLM reasoning in an explicit and queryable metric world representation. By separating spatial state estimation from semantic planning and adding counterfactual branch reasoning, GTA produces physically grounded waypoints without task-specific fine-tuning. Experiments on R2R-CE and RxR-CE show state-of-the-art zero-shot performance, while real-world deployment demonstrates that the same representation supports transfer across robot embodiments. These findings suggest that a key bottleneck in MLLM-based embodied navigation lies not only in the reasoning model itself, but also in how the physical world is represented. GTA inherits the assumptions of the proposed world representation. It relies on reasonably accurate depth, pose, and semantic perception, and currently focuses on mostly static scenes. Performance may therefore degrade under severe localization drift, perception failures, or crowded dynamic environments. Extending the representation with open-vocabulary 3D perception and dynamic scene modeling is left for future work. The bounded history window also motivates sparse memory retrieval for longer-horizon applications.

References

[1] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. Van Den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3674–3683, 2018.

[2] G. Comanici, E. Bieber, M. Schaekermann, I. Pasupat, N. Sachdeva, I. Dhillon, M. Blistein, O. Ram, D. Zhang, E. Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.

[3] S. Bai, Y. Cai, R. Chen, K. Chen, X. Chen, Z. Cheng, L. Deng, W. Ding, C. Gao, C. Ge, W. Ge, Z. Guo, Q. Huang, J. Huang, F. Huang, B. Hui, S. Jiang, Z. Li, M. Li, M. Li, K. Li, Z. Lin, J. Lin, X. Liu, J. Liu, C. Liu, Y. Liu, D. Liu, S. Liu, D. Lu, R. Luo, C. Lv, R. Men, L. Meng, X. Ren, X. Ren, S. Song, Y. Sun, J. Tang, J. Tu, J. Wan, P. Wang, P. Wang, Q. Wang, Y. Wang, T. Xie, Y. Xu, H. Xu, J. Xu, Z. Yang, M. Yang, J. Yang, A. Yang, B. Yu, F. Zhang, H. Zhang, X. Zhang, B. Zheng, H. Zhong, J. Zhou, F. Zhou, J. Zhou, Y. Zhu, and K. Zhu. Qwen3-vl technical report, 2025. URL https://arxiv.org/abs/2511.21631.

[4] G. Zhou, Y. Hong, and Q. Wu. Navgpt: Explicit reasoning in vision-and-language navigation with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 7641–7649, 2024.

[5] G. Zhou, Y. Hong, Z. Wang, X. E. Wang, and Q. Wu. Navgpt-2: Unleashing navigational reasoning capability for large vision-language models. In *European Conference on Computer Vision*, pages 260–278. Springer, 2024.

[6] P. Chen, X. Sun, H. Zhi, R. Zeng, T. H. Li, G. Liu, M. Tan, and C. Gan. A2nav: Action-aware zero-shot robot navigation by exploiting vision-and-language ability of foundation models. *arXiv preprint arXiv:2308.07997*, 2023.

[7] Y. Qiao, W. Lyu, H. Wang, Z. Wang, Z. Li, Y. Zhang, M. Tan, and Q. Wu. Open-nav: Exploring zero-shot vision-and-language navigation in continuous environment with open-source llms. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6710–6717. IEEE, 2025.

[8] J. Chen, B. Lin, R. Xu, Z. Chai, X. Liang, and K.-Y. Wong. Mapgpt: Map-guided prompting with adaptive path planning for vision-and-language navigation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9796–9810, 2024.

[9] L. Zhang, X. Hao, Q. Xu, Q. Zhang, X. Zhang, P. Wang, J. Zhang, Z. Wang, S. Zhang, and R. Xu. MapNav: A novel memory representation via annotated semantic maps for VLM-based vision-and-language navigation. In W. Che, J. Nabende, E. Shutova, and M. T. Pilehvar, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13032–13056, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0.

[10] H. Wang, W. Wang, W. Liang, C. Xiong, and J. Shen. Structured scene memory for vision-language navigation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 8455–8464, 2021.

[11] Z. Wang, X. Li, J. Yang, Y. Liu, and S. Jiang. Gridmm: Grid memory map for vision-and-language navigation. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 15625–15636, 2023.

[12] X. Shi, Z. Li, W. Lyu, J. Xia, F. Dayoub, Y. Qiao, and Q. Wu. Smartway: Enhanced waypoint prediction and backtracking for zero-shot vision-and-language navigation. *arXiv preprint arXiv:2503.10069*, 2025.

[13] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *European Conference on Computer Vision*, pages 104–120. Springer, 2020.

[14] Y. Qi, Q. Wu, P. Anderson, X. Wang, W. Y. Wang, C. Shen, and A. v. d. Hengel. Reverie: Remote embodied visual referring expression in real indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9982–9991, 2020.

[15] J. Zhu, W. Wang, Z. Chen, Z. Liu, S. Ye, L. Gu, H. Tian, Y. Duan, W. Su, J. Shao, et al. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

[16] OpenAI. Gpt-5 system card. https://cdn.openai.com/gpt-5-system-card.pdf, 2025.

[17] 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.

[18] D. Fried, R. Hu, V. Cirik, A. Rohrbach, J. Andreas, L.-P. Morency, T. Berg-Kirkpatrick, K. Saenko, D. Klein, and T. Darrell. Speaker-follower models for vision-and-language navigation. *Advances in neural information processing systems*, 31, 2018.

[19] M. Z. Irshad, C.-Y. Ma, and Z. Kira. Hierarchical Cross-Modal Agent for Robotics Vision-and-Language Navigation, Apr. 2021. URL http://arxiv.org/abs/2104.10674. arXiv:2104.10674 [cs].

[20] M. Z. Irshad, N. C. Mithun, Z. Seymour, H.-P. Chiu, S. Samarasekera, and R. Kumar. SASRA: Semantically-aware Spatio-temporal Reasoning Agent for Vision-and-Language Navigation in Continuous Environments, Aug. 2021. URL http://arxiv.org/abs/2108.11945. arXiv:2108.11945 [cs].

[21] S. Raychaudhuri, S. Wani, S. Patel, U. Jain, and A. X. Chang. Language-Aligned Waypoint (LAW) Supervision for Vision-and-Language Navigation in Continuous Environments, Sept. 2021. URL http://arxiv.org/abs/2109.15207. arXiv:2109.15207 [cs].

[22] Y. Hong, Z. Wang, Q. Wu, and S. Gould. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15439–15449, 2022.

[23] J. Krantz, A. Gokaslan, D. Batra, S. Lee, and O. Maksymets. Waypoint models for instruction-guided navigation in continuous environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15162–15171, 2021.

[24] Z. Li, G. Zhou, H. Hong, Y. Shao, W. Lyu, Y. Qiao, and Q. Wu. Ground-level viewpoint vision-and-language navigation in continuous environments. *IEEE International Conference on Robotics and Automation*, 2025.

[25] D. An, H. Wang, W. Wang, Z. Wang, Y. Huang, K. He, and L. Wang. Etpnav: Evolving topological planning for vision-language navigation in continuous environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

[26] D. An, Y. Qi, Y. Li, Y. Huang, L. Wang, T. Tan, and J. Shao. Bevbert: Multimodal map pre-training for language-guided navigation. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.

[27] J. Zhang, K. Wang, R. Xu, G. Zhou, Y. Hong, X. Fang, Q. Wu, Z. Zhang, and H. Wang. Navid: Video-based vlm plans the next step for vision-and-language navigation. *CoRR*, 2024.

[28] J. Zhang, K. Wang, S. Wang, M. Li, H. Liu, S. Wei, Z. Wang, Z. Zhang, and H. Wang. Uni-navid: A video-based vision-language-action model for unifying embodied navigation tasks. *arXiv preprint arXiv:2412.06224*, 2024.

[29] J. Zhang, A. Li, Y. Qi, M. Li, J. Liu, S. Wang, H. Liu, G. Zhou, Y. Wu, X. Li, et al. Embodied navigation foundation model. *arXiv preprint arXiv:2509.12129*, 2025.

[30] Y. Long, X. Li, W. Cai, and H. Dong. Discuss before moving: Visual language navigation via multi-expert discussions. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 17380–17387. IEEE, 2024.

[31] J. Chen, B. Lin, X. Liu, L. Ma, X. Liang, and K.-Y. K. Wong. Affordances-oriented planning using foundation models for continuous vision-language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23568–23576, 2025.

[32] Y. Long, W. Cai, H. Wang, G. Zhan, and H. Dong. Instructnav: Zero-shot system for generic instruction navigation in unexplored environment. In *Conference on Robot Learning*, pages 2049–2060. PMLR, 2025.

[33] K. Chen, D. An, Y. Huang, R. Xu, Y. Su, Y. Ling, I. Reid, and L. Wang. Constraint-aware zero-shot vision-language navigation in continuous environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

[34] B. Li, S. Li, W. Wang, A. Li, Z. Cao, and H. X. Liu. Boosting zero-shot vln via abstract obstacle map-based waypoint prediction with topograph-and-visitinfo-aware prompting. *arXiv preprint arXiv:2509.20499*, 2025.

[35] N. Hughes, Y. Chang, and L. Carlone. Hydra: A real-time spatial perception system for 3d scene graph construction and optimization. *arXiv preprint arXiv:2201.13360*, 2022.

[36] K. Rana, J. Haviland, S. Garg, J. Abou-Chakra, I. Reid, and N. Suenderhauf. Sayplan: Grounding large language models using 3d scene graphs for scalable robot task planning. *arXiv preprint arXiv:2307.06135*, 2023.

[37] A. Ray, J. Arkin, H. Biggie, C. Fan, L. Carlone, and N. Roy. Structured interfaces for automated reasoning with 3d scene graphs. In *2026 IEEE International Conference on Robotics and Automation (ICRA)*, pages 18418–18425. IEEE, 2026.

[38] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.

[39] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*, pages 127–136. Ieee, 2011.

[40] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

[41] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115, 2023.

[42] Y. Hong, Q. Wu, Y. Qi, C. Rodriguez-Opazo, and S. Gould. Vln bert: A recurrent vision-and-language bert for navigation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 1643–1653, 2021.

[43] Z. Wang, X. Li, J. Yang, Y. Liu, J. Hu, M. Jiang, and S. Jiang. Lookahead exploration with neural radiance representation for continuous vision-language navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13753–13762, 2024.

[44] D. Zheng, S. Huang, Y. Li, and L. Wang. Efficient-vln: A training-efficient vision-language navigation model. *arXiv preprint arXiv:2512.10310*, 2025.

[45] D. He, X. Gao, H. Li, J. Han, and D. Zhang. Strider: Navigation via instruction-aligned structural decision space optimization. *arXiv preprint arXiv:2511.00033*, 2025.

[46] N. P. Bhatt, Y. Yang, R. Siva, P. Samineni, D. Milan, Z. Wang, and U. Topcu. Vln-zero: Rapid exploration and cache-enabled neurosymbolic vision-language planning for zero-shot transfer in robot navigation. *arXiv preprint arXiv:2509.18592*, 2025.

[47] A. Ku, P. Anderson, R. Patel, E. Ie, and J. Baldridge. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4392–4412, 2020.

[48] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9339–9347, 2019.

[49] J. Zhang, Z. Li, S. Wang, X. Shi, Z. Wei, and Q. Wu. Spatialnav: Leveraging spatial scene graphs for zero-shot vision-and-language navigation. *arXiv preprint arXiv:2601.06806*, 2026.

[50] L. Wang, X. Xia, H. Zhao, H. Wang, T. Wang, Y. Chen, C. Liu, Q. Chen, and J. Pang. Rethinking the embodied gap in vision-and-language navigation: A holistic study of physical and visual disparities. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9455–9465, 2025.

[51] Z. Wang, J. Li, Y. Hong, Y. Wang, Q. Wu, M. Bansal, S. Gould, H. Tan, and Y. Qiao. Scaling data generation in vision-and-language navigation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12009–12020, 2023.

A Experimental Setup Details

Benchmarks. The main continuous-navigation evaluations use R2R-CE and RxR-CE in Habitat [48]. The world-representation comparison in Fig. 4 additionally includes discrete R2R and REVERIE evaluations. The main R2R-CE result uses the full val-unseen split of 1,839 episodes. The main RxR-CE result uses a scene-stratified sample of 260 English val-unseen episodes with seed 1024. The 180-episode R2R-CE ablation set filters trivial cases and prioritizes long-horizon, multi-room trajectories. Additional R2R-CE comparisons and mechanism ablations use Open-Nav's released 100-episode set [7, 12, 45, 49]. The additional RxR-CE evaluation uses a random sample of 300 English val-unseen episodes. BZS-VLN's episode IDs were unavailable, so this comparison does not establish episode-level matching.

Metrics. We report Success Rate (SR), Oracle Success Rate (OSR), Success weighted by Path Length (SPL), Navigation Error (NE), and normalized Dynamic Time Warping (nDTW), the last assessing trajectory fidelity to the ground-truth path. An episode succeeds if the agent stops within 3.0 m of the target.

Baselines. We compare against two groups (Tab. 1): supervised agents (#1–#8) trained on in-domain data [22, 11, 25, 26, 43, 29, 44], from CMA and VLN-BERT to the recent state-of-the-art Efficient-VLN, and training-free zero-shot agents (#9–#14) [7, 33, 12, 45, 46, 34], of which BZS-VLN (#14) is the prior best.

B GTA Implementation Details

The volumetric map $\mathcal{M}_{vol}$ uses a TSDF whose confidence weight ω_t decays with depth and grazing angle, with observations fused via the running-average update (Eq. (2)). The topological graph $\mathcal{G}_{topo}$ merges a new position into the nearest node when their distance is below δ_{merge}, and otherwise instantiates a new node. A loop alert is raised when the same node is visited at least three times within a six-step window, or when the agent oscillates between two nodes. At each step we select four egocentric views at the canonical angles $\{0°, 90°, 180°, 270°\}$ from the rotating RGB-D camera, and overlay an integer coordinate grid over $[0, 1000]^2$ on the BEV and egocentric images. The zero-shot reasoner is GPT-5.1, accessed via its official API. For ablations, the 180-episode R2R-CE subset (Sec. 4.1) is sampled from the val-unseen split, favoring episodes with long reference trajectories and complex, fine-grained instructions. Table 7 lists the full hyperparameters of the zero-shot system, and Algorithm 1 summarizes the closed-loop interaction.

Table 7: GTA zero-shot hyperparameters.

Component	Setting
TSDF voxel size	0.05 m
TSDF truncation μ	5×vox
Volume bounds	$[-15, 15]^2$
Merge thr. δ_{merge}	0.8 m
Loop thr. c_{loop}	3
History window H	20
Floor thr. $\lvert\Delta h\rvert$	0.3 m
Camera / FOV	256^2/90°
View angles	$\{0, 90, 180, 270\}^\circ$
Coordinate grid	$[0, 1000]^2$
Step (norm/near)	0.5/0.25 m
Success thr.	3.0 m
Max steps	50
Reasoner	GPT-5.1

C GTA-Light Training Details

Training-data construction. Data are generated from the same world representation as the zero-shot system, using scenes from the R2R training split. No validation scene, episode, or instruction is used for this data construction or student training. For each scene, we sample a start $\mathbf{p}_0$ and target $\mathbf{p}_{goal}$ and compute a reference trajectory by $A^\star$ search on the navigable graph. Instructions paired with these trajectories are generated using ScaleVLN [51]. The target is used only for offline data generation; the deployed agent never accesses $\mathbf{p}_{goal}$ and stops when $\mathcal{F}$ emits a stop action. We replay each trajectory while building $\mathcal{W}_t$ exactly as in the online setting and query the teacher $\mathcal{P}_t^\star = \mathcal{F}_{tea}(\mathrm{Render}(\mathcal{W}_t), \mathcal{I})$ to obtain the supervised pairs $\mathcal{D} = \{(\mathcal{W}_t, \mathcal{P}_t^\star)\}$. This turns the expensive interpreter into an offline data engine and amortizes its reasoning into a compact model. Training takes approximately 70 hours on four H100 GPUs. Table 8 summarizes the training configuration.

Robustness augmentation. To stay stable under the noisy observations encountered in real navigation, we perturb both the input state and the target prompt during training. The input-state perturbations $\mathcal{A}_\epsilon^W$ comprise pose/heading jitter, dropped graph nodes, masked TSDF cells, depth noise, and varied render style. The prompt perturbations $\mathcal{A}_\eta^P$ comprise reordered non-causal fields, paraphrased alerts, multi-resolution coordinate quantization, and alternative valid waypoints along the same A* corridor. These expose the model to localization drift, partial observability, and rendering shifts.

Algorithm 1 Overview of the GTA Pipeline.

Require: Global instruction $\mathcal{I}$, start pose $\mathbf{p}_0$, frozen MLLM $\mathcal{F}$
Require: Thresholds: δ_{merge} (node merging), c_{loop} (loop detection), T_{max} (max steps)
Initialize: $\mathcal{M}_{vol} \leftarrow \emptyset$, $\mathcal{G}_{topo} \leftarrow (\{\mathbf{p}_0\}, \emptyset)$, $\mathcal{H} \leftarrow \emptyset$, checkpoint stack $\mathcal{K} \leftarrow \emptyset$
Initialize checklist: $\mathcal{C}_t \leftarrow$ DECOMPOSE($\mathcal{I}$)
while $\mathcal{F}$ has not emitted STOP **and** $t < T_{max}$ **do**
  // 1. Metric perception (high-frequency)
  Capture rotating view $\mathcal{O}_s = (\mathbf{I}_s, \mathbf{D}_s, \theta_s)$; update $\mathcal{M}_{vol}$ via TSDF (Eq. (2))
  // 2. Topological memory & safety
  $v_{curr} \leftarrow$ CLUSTERNODE($\mathcal{G}_{topo}, \mathbf{p}_t, \delta_{merge}$); $\mathcal{S}_{alert} \leftarrow \emptyset$
  if $c_{v_{curr}} > c_{loop}$ **then**
    $\mathcal{S}_{alert} \leftarrow$ "Potential loop detected"
  end if
  // 3. Verify previous branch against real observation (Eq. (8))
  if $\mathcal{K} \neq \emptyset$ **and** $\mathcal{F}$ judges observation inconsistent with stored $e_{t'}^{(k^\star)}$, **or** $\mathcal{S}_{alert} \neq \emptyset$ **then**
    Pop to latest $\kappa_{t'} \in \mathcal{K}$; block failed branch: $\mathcal{B}_{t'}^{\text{blk}} \leftarrow \mathcal{B}_{t'}^{\text{blk}} \cup \{k^\star\}$
    Physically retrace visited nodes in $\mathcal{G}_{topo}$ to $\mathbf{p}_{t'}$ using π_{local}
    if $\mathcal{B}_{t'} \setminus \mathcal{B}_{t'}^{\text{blk}} = \emptyset$ **then** release $\kappa_{t'}$, inject backtrack alert into $\mathcal{S}_{alert}$
  end if
  // 4. Interface: serialize $\mathcal{W}_t$ into prompt $\mathcal{P}_t$
  $\{\mathbf{I}_{ego}^{(k)}\}_{k=1}^4 \leftarrow$ SELECTVIEWS($\mathcal{N}_t$) (Eq. (4))
  $\mathcal{V}_{prompt} \leftarrow [\mathbf{I}_{bev}, \{\mathbf{I}_{ego}^{(k)}\}_{k=1}^4]$
  $\mathcal{T}_{prompt} \leftarrow \mathcal{C}_t \oplus$ STATE($\mathcal{G}_{topo}, \mathcal{S}_{alert}$) $\oplus$ HISTORY($\mathcal{H}$) $\oplus \mathcal{I}$
  $\mathcal{P}_t \leftarrow \{\mathcal{V}_{prompt}, \mathcal{T}_{prompt}\}$
  // 5. Reasoning: act, and checkpoint at branch points
  $\mathcal{R} \leftarrow \mathcal{F}(\mathcal{P}_t)$; parse (m_t, u_t, v_t), $\mathcal{C}_{t+1} \leftarrow$ PARSEJSON($\mathcal{R}$)
  $\mathcal{B}_t \leftarrow$ directions marked by $\mathcal{F}$ as traversable and relevant to $\mathcal{C}_t$ (from $\mathcal{R}$)
  if $|\mathcal{B}_t| > 1$ **then**
    $\mathcal{F}$ emits per-direction expectations $\{e_t^{(k)}\}_{k \in \mathcal{B}_t}$
    Push $\kappa_t = \langle \mathbf{p}_t, \{(k, e_t^{(k)})\}_{k \in \mathcal{B}_t}, \emptyset \rangle$ onto $\mathcal{K}$
  end if
  // 6. Grounding & execution
  $\mathbf{w}_t \leftarrow$ RAYCAST($\mathbf{p}_{cam}^{(m_t)}, \mathbf{r}^{(m_t)}(u_t, v_t), \mathcal{M}_{vol}$) (Eq. (9))
  $\mathbf{c}_t \leftarrow \pi_{local}(\mathbf{p}_t, \mathbf{w}_t)$
  Execute $\mathbf{c}_t$, advance to $\mathbf{p}_{t+1}$, update history $\mathcal{H}$
end while

D Real-World Deployment Details

Wheeled robot. For ground-based tasks we use a TurtleBot 4 mobile base equipped with an Intel RealSense D455 RGB-D camera. The camera is mounted on a servo motor to perform the orthogonal panoramic scanning (4 views) described in Sec. 3. Low-level motion planning is handled by the ROS 2 Nav2 stack, while MLLM inference runs on a remote server communicating via standard HTTP requests.

Table 8: GTA-Light training hyperparameters.

Component	Setting
Student backbone	InternVL3-8B
Teacher	GPT-5.1
LoRA rank r / α	16 / 32
Trainable modules	LoRA adapters + multimodal projector
Quantization	4-bit NF4 (QLoRA)
Learning rate	2×10^{-4}
Batch size	128 effective batch size
Epochs	3
Optimizer	paged AdamW 8-bit
Loss weights λ_{view} / λ_{coord}	1.0 / 2.0
Training trajectories / pairs	10k / 180k

Quadruped robot. For legged locomotion we use a Unitree Go2 quadruped equipped with an Intel RealSense RGB-D camera, mounted on a servo to perform the same orthogonal four-view scanning as the wheeled platform. We deliberately reuse the wheeled robot's perception and navigation setup unchanged, including the ROS 2 motion stack and the same remote MLLM-inference interface, so that the only variable across the two ground platforms is the locomotion dynamics. This isolates embodiment as the sole factor and tests whether the world representation remains a valid interface under the irregular, non-planar motion of a legged body.

Aerial robot. For aerial operations we employ a custom-built quadrotor integrated with an Intel RealSense D435 RGB-D camera. The UAV uses a Betaflight-based flight controller, while high-level command computation and communication run on a CoolPi 4B onboard computer. Precise state estimation and global localization are provided by an external motion-capture system.

Protocol. We conduct 50 navigation trials with complex language instructions on each platform in unseen real-world environments. The three embodiments span a wide range of motion dynamics and camera viewpoints, from the stable planar motion of the wheeled base, through the irregular gait-induced jitter of the quadruped, to the free six-DoF motion and elevated viewpoint of the aerial platform, while sharing the same GTA framework without platform-specific tuning.

E Wall-Clock Latency and Real-Time Viability

Table 9: Inference latency per reasoning step, measured over 100 trials.

Module / Backbone	Mean Latency (s)	Std. Dev. (s)
Perception & Mapping		
TSDF Integration + BEV Rendering	0.0972	$\pm$0.0844
Reasoning		
GTA-Light (Local)	3.96	$\pm$0.47
Gemini-2.5-Flash (API)	3.27	$\pm$0.37
GPT-5.1 (API)	5.42	$\pm$1.92

A practical concern for any MLLM-based navigation system is whether the reasoning latency is compatible with real-time deployment. We address this through two complementary arguments: (1) an asynchronous execution paradigm that structurally decouples perception from reasoning, and (2) empirical latency measurements across all backbone configurations.

Asynchronous Execution Paradigm. GTA adopts a hierarchical, asynchronous design in which the *Perception & Mapping* module runs at high frequency while the *Counterfactual Reasoning Brain* operates as a low-frequency "Slow System." Rather than blocking navigation on each MLLM call, the system initiates the next reasoning query as soon as the robot departs toward its current waypoint. At a safe indoor cruising speed of 0.1 m/s, the robot requires approximately 5–10 s to traverse a typical inter-waypoint distance of 0.5–1.0 m. The measured mean reasoning latencies support

overlapping inference with motion on typical segments, although latency variability can still cause pauses on shorter segments.

Latency Measurements. Table 9 reports the mean inference latency and standard deviation per reasoning step, measured over 100 trials for each module and backbone. TSDF integration and BEV rendering together take 0.0972 s on average, contributing little overhead. Mean reasoning latency ranges from 3.27 s (Gemini-2.5-Flash) to 5.42 s (GPT-5.1), supporting overlap with motion on typical segments.

F Additional Evaluation Details

Component ablations. On Open-Nav's 100-episode R2R-CE set, Tab. 10 reports SR changes in percentage points relative to full GTA (47.0% SR). Removing egocentric views causes the largest drop, followed by BEV and topology. Removing the coordinate grid leaves aggregate SR unchanged on this set; this does not establish its redundancy across models or environments.

Table 10: Component ablations

Ablated component	ΔSR (pp)
BEV	-5.0
Topological graph	-3.0
Egocentric views	-16.0
Coordinate grid	0.0

G Notation

Table 11: Main notation used in GTA.

Symbol	Meaning
$\mathcal{I}$	Global navigation instruction.
$\mathcal{F}$, ϕ, ϕ_θ	Frozen MLLM reasoner, training-free interface, and learned interpreter.
$\mathbf{s}_t$, $\mathbf{p}_t$	Agent pose (x_t, y_t, θ_t) and planar position.
$a_t = (m_t, u_t, v_t)$	Action comprising a view index and target pixel.
$\mathcal{W}_t = \langle \mathcal{M}_{vol}, \mathcal{G}_{topo} \rangle$	World representation: TSDF volume and topological graph.
$S_t(\mathbf{v})$, $\Omega_t(\mathbf{v})$	Fused signed distance and accumulated voxel weight.
ψ_t, ω_t, μ	Per-observation signed distance, fusion weight, and truncation radius.
$\mathbf{T}_{wb}$, $\mathbf{T}_{bc}$, $\mathbf{K}$	Body pose, camera extrinsics, and camera intrinsics.
$v_i = \langle \mathbf{p}_i, c_i \rangle$	Topological node storing position and visit count.
δ_{merge}, c_{loop}	Node-merging distance and loop-detection visit threshold.
$\mathcal{N}_t$, $\mathbf{q}_s$, $\mathbf{q}_t$	Candidate past frames, frame orientation, and current agent orientation.
$\mathbf{I}_{bev}$, $\mathbf{I}_{ego}^{(k)}$	Rendered map view and selected egocentric view.
$\mathcal{P}_t$, $\mathcal{V}_{prompt}$, $\mathcal{T}_{prompt}$	Multimodal prompt and its visual and textual components.
$\mathcal{C}_t$, $\mathcal{S}_{alert}$, $\mathcal{H}_t$, H	Dynamic checklist, safety alerts, execution history, and history-window length.
$\mathcal{B}_t$, $\mathcal{B}_t^{\text{blk}}$	Viable branches and blocked directions.
$e_t^{(k)}$, $k^\star$, $g(\cdot)$	Branch expectation, selected branch, and the MLLM's preference over remaining branches.
κ_t, $\mathcal{K}$	Checkpoint and checkpoint stack.
$\mathbf{w}_t$, π_{local}	Grounded metric waypoint and local planner.

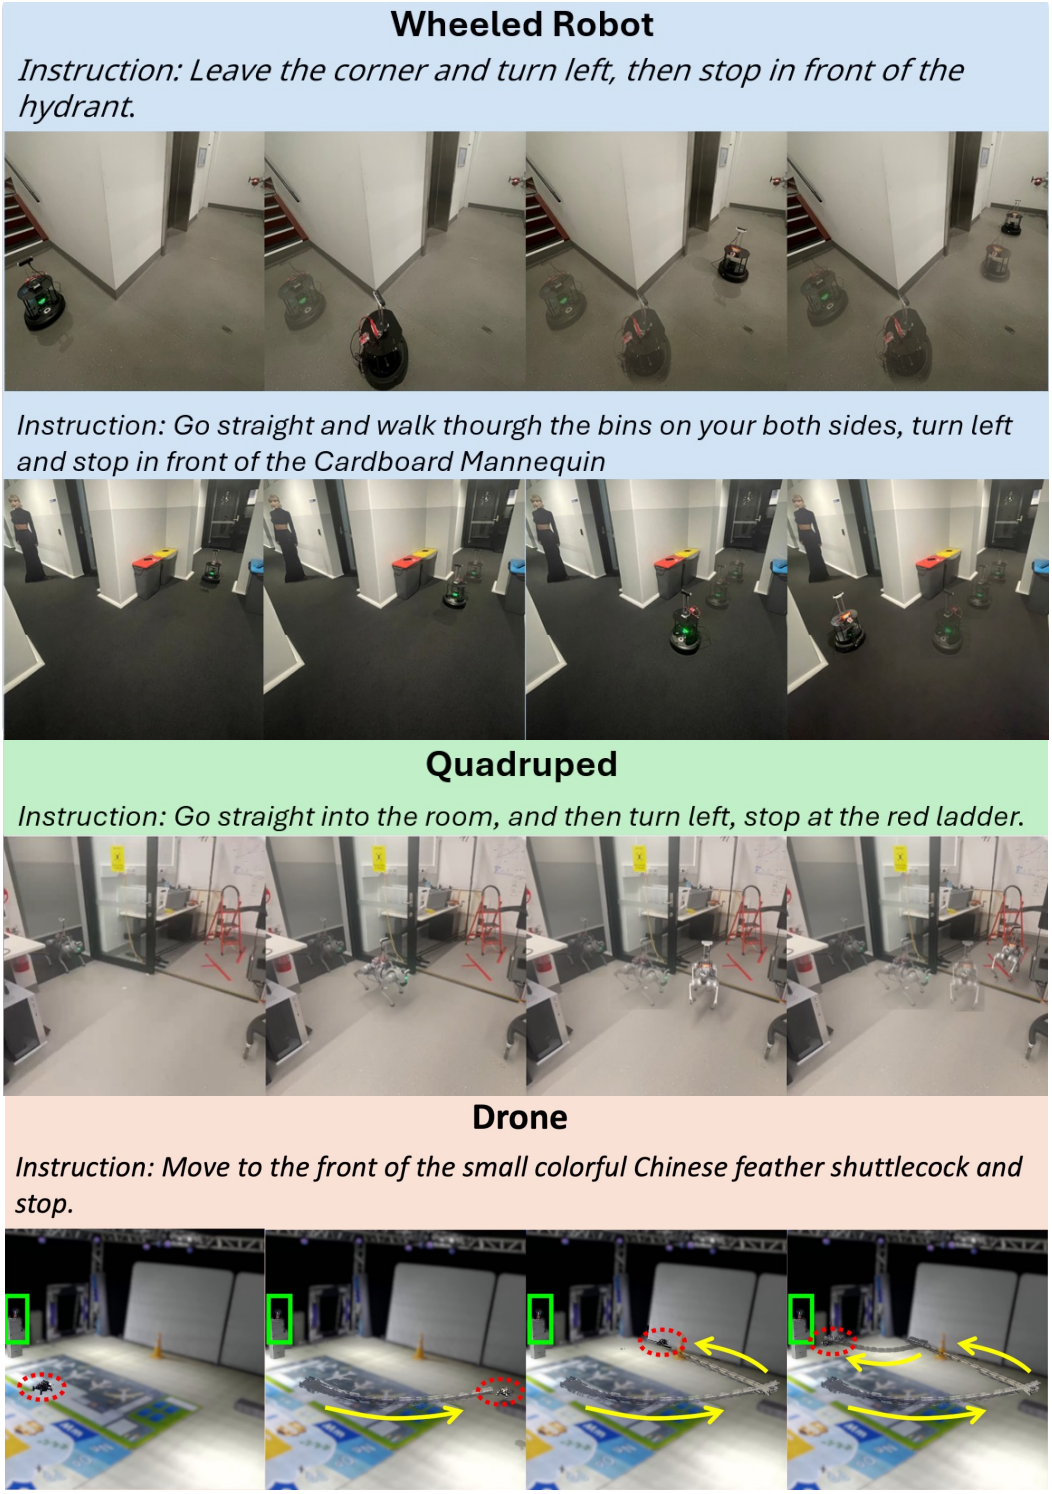


Figure 5: **Zero-shot sim-to-real transfer across diverse robot embodiments.** We demonstrate the generalization of GTA by deploying it on distinct physical platforms in unseen real-world environments.