

GPT-6-Astra in a Navigation Workflow

Behavioral Analysis in Zero-Shot Vision-and-Language Navigation in Continuous Environments

Guangzhao Dai¹, Qi Wu², Bin Zhu^{1†}

¹School of Computing and Information Systems, Singapore Management University

²Australia Institute for Machine Learning

[†]Corresponding author and project lead

Website: <https://daiguangzhao.github.io/gpt-6-astra-vln/>




Figure 1. **From interpreting instructions to completing navigation.** Within the evaluated workflow, GPT-6-Astra can match landmarks and earlier actions to instructions, seek additional evidence, and revise uncertain judgments. These judgments do not always lead to progress or appropriate stopping. The five perspectives (Sections 5.1–5.5) organize this behavioral analysis; the mountain is a conceptual illustration.

Abstract

We study GPT-6-Astra in a zero-shot Vision-and-Language Navigation in Continuous Environments (VLN-CE) system, where it interprets instructions, assesses its surroundings, and proposes actions. The system uses a common observation–decision–execution workflow with direct model API calls, without a packaged agent harness or navigation-specific fine-tuning. In this workflow, each request receives selected observations, execution feedback, and retained progress records. Evaluation covers the complete system, including context management and action control. We evaluate the system on 50 of the 100 R2R-CE val-unseen episodes used by Open-Nav. It achieves a

success rate of 52.0%, an SPL of 48.9%, and an nDTW of 70.8%. Our analysis highlights three findings. First, recorded responses link landmarks and earlier actions to instructions using observations and supplied history. Second, reviews include requests for additional views and revisions of uncertain judgments. Third, the results suggest a gap between task understanding and autonomous completion: an unfinished crossing is recognized while rotation continues. At termination, 36.0% of episodes succeed with a workflow-accepted STOP, while another 16.0% meet the distance criterion at the step limit. These results highlight a central challenge: translating correct local judgments into sustained progress and appropriate stopping.

1. Introduction

We study how GPT-6-Astra uses visual observations and navigation history to decide where to go, what to inspect, and when to stop. Its recorded responses contain detailed judgments about which landmark matches an instruction, what earlier actions accomplished, and what evidence is still missing. Such judgments are useful when several targets look plausible or the effect of a previous action is unclear. Their value for navigation, however, depends on whether they help the agent choose and complete the next action. This motivates following the model’s judgments through an entire navigation task.

Vision-and-Language Navigation (VLN) provides a setting for this analysis: an agent must follow a natural-language instruction through a visual environment [3]. In continuous environments (VLN-CE), it must execute turns and translations to follow the route and then stop at the destination [13]. Recognizing an entrance does not establish that the agent has crossed it, and approaching a goal does not establish that it will stop appropriately. We therefore ask: *can GPT-6-Astra use the supplied observations and action history to guide an agent to the instructed destination and determine when to stop?*

Prior zero-shot navigation systems combine language-model decisions with visual observations, maps, history, and feedback [4, 19–21, 38]. Recent work further explores trajectory imagination, targeted perception, and tool-based verification [15, 24, 31]. These approaches illustrate how observations and external mechanisms support navigation decisions. NavBench also distinguishes navigation comprehension from execution [18]. Building on this distinction, we examine what GPT-6-Astra judges and what the navigation system subsequently executes within the same episode.

We use an independently implemented, common observation–decision–execution workflow and evaluate the subset listed in Table 1. GPT-6-Astra interprets observations, assesses progress, proposes actions, and reviews arrival evidence through direct model API calls, without a packaged agent harness or navigation-specific fine-tuning. External code selects the context for each request, retains prior observations and judgments, schedules reviews, and constrains execution. We therefore report the performance of this complete system and interpret the model’s responses in light of the evidence and control mechanisms available at each step.

Our analysis links saved observations and model responses to executed actions, trajectories, and completion outcomes. The five perspectives in Figure 1 organize this evidence and support three findings, each connecting a model judgment to its observed consequences.

(1) GPT-6-Astra shows strengths in linking landmarks and earlier actions to instructions. For example, it distinguishes the “second on the left” from neighboring openings

and uses later views to confirm that an earlier turn matched the route instruction. The latter response preserves when the turn occurred while clarifying what it accomplished. Across the evaluated episodes, the timing and order analyses each find supported evidence for every applicable check in 31 episodes (Section 5.3). These judgments use supplied history. They show that later evidence can establish an instruction match for an action that has already occurred, so the time of an event and the time of its confirmation need to be distinguished.

(2) GPT-6-Astra shows strengths in seeking visual information and revising uncertain judgments. It proposes moving past a door leaf, after which the executed movement reveals a hidden passage. In another case, additional views expose shelving and a back wall, leading it to reject an uncertain hallway candidate. Our analysis finds that 36 of 66 adjustment processes help overcome a local difficulty, while 48 of 75 verification processes resolve an uncertain judgment with supporting evidence (Sections 5.4–5.5). These are system processes, with reviews scheduled by the workflow. They demonstrate the value of obtaining evidence that answers a specific question, while leaving open whether resolving that question leads to further route progress.

(3) GPT-6-Astra shows limitations in translating task understanding into autonomous completion within this workflow. In a stalled doorway sequence, the model correctly distinguishes a completed turn from a pending crossing, yet the system continues rotating without crossing. A gap between physical arrival and accepted completion also appears in the aggregate results. The system achieves 52.0% SR, 48.9% SPL, and 70.8% nDTW, but only 36.0% of episodes succeed with a workflow-accepted STOP; another 16.0% meet the distance criterion at the step limit. Thus, supported judgments and successful endpoints do not by themselves establish effective task completion. Assessing this system requires checking whether the required motion follows a judgment and whether arrival leads to an appropriate stopping decision.

2. Task and Navigation Workflow

Task formulation. Vision-and-Language Navigation in Continuous Environments (VLN-CE) requires an agent to follow a natural-language instruction $\mathcal{I}$ and reach the instructed destination through low-level actions [13]. In the evaluated setting, the action space $\mathcal{A}$ comprises forward movement of 0.25 m, left/right rotation of 15°, and STOP. The agent receives an initial panorama for orientation, followed by egocentric RGB/depth observations $\mathcal{O}_t$. Selected historical observations, executed actions, and execution feedback form the available history $\mathcal{H}_t$.

The navigation decision is

$$a_t = \pi_{\text{agent}}(\mathcal{I}, \mathcal{O}_t, \mathcal{H}_t), \quad a_t \in \mathcal{A}, \quad (1)$$




Figure 2. **Model calls within an externally managed navigation loop.** (a) Observations supply scene evidence. (b) External state retains history and prior judgments; the workflow selects context and constructs separate calls for the four model roles. Outputs update retained records. (c) Action constraints and STOP checks precede execution, whose feedback updates the next request.

where π_{agent} denotes the complete navigation policy, combining GPT-6-Astra’s proposals with external context management and action constraints. Executing a_t produces the next observation and motion/collision feedback, which inform subsequent decisions. The GT route and endpoint are reserved for offline evaluation.

Following the instruction requires identifying the relevant landmarks, tracking completed actions, and obtaining more evidence when the next move is unclear. We examine whether these judgments lead to the required movement and an appropriate stopping decision.

Division of responsibilities. Figure 2 shows our implementation of the common observation–decision–execution loop using direct model API calls, without a packaged agent harness. The instruction and an initial scan of 12 headings establish the starting context; subsequent observations are frontal RGB-D. GPT-6-Astra serves four prompted roles: task planning, progress assessment, action proposal, and arrival assessment. Workflow code schedules these calls, selects their evidence, validates and retains structured outputs, and applies movement and completion constraints. The model can request additional saved images or propose a new viewpoint; the workflow supplies the selected images or executes an action after applying its constraints. The loop ends at a workflow-accepted STOP or the episode budget. Context selection and action control therefore contribute to the reported system performance alongside the model’s judgments.

3. Evaluation Setup

3.1. Dataset

The evaluated subset comprises the first 50 episodes in ascending episode-ID order from the Open-Nav evaluation pool [19], which is also used by SmartWay and Fast-SmartWay [20, 21]. It spans nine reconstructed indoor scenes from the unseen-environment split. The evaluated episodes use the R2R_VLNCE_v1-3 release, with corrected initial headings and matching GT trajectories [22]. Table 1 lists all evaluated episode IDs and their scene distribution.

3.2. Context Management and Action Interface

Each generation is a separate chat request containing a role-specific system prompt and one newly constructed multimodal user message. The implementation does not replay the full assistant-message history or pass reasoning-state continuation items between requests. Earlier observations, execution receipts, task progress, and selected model judgments are retained externally and supplied in later requests. Decisions can therefore use navigation history, with the workflow determining which parts of that history remain available.

Ordinary action calls receive up to four selected observation frames and an eight-action receipt window, with the task-activation view retained when needed. Progress reviews can draw on the episode archive, event records, prior questions, and validation feedback; arrival reviews additionally receive the two most recent review records.

Table 1. **Composition of the evaluated subset.** All 50 episodes are grouped by scene; execution follows ascending episode-ID order across the subset.

Scene ID	#Ep.	Episode IDs
2azQ1b91cZZ	9	11, 70, 116, 140, 166, 265, 330, 403, 513
EU6Fwq7SyZv	3	321, 348, 461
QUCTc6BB5sX	7	94, 150, 176, 190, 191, 423, 454
TbHJrupSAjP	8	171, 181, 259, 275, 308, 338, 377, 411
X7HyMhZNoso	4	218, 226, 244, 371
Z6MFQCViBuw	3	207, 432, 469
oLBMNvg9in8	2	232, 479
x8F5xyUWy9e	3	7, 187, 312
zsNo4HB9uLZ	11	13, 40, 42, 52, 156, 247, 362, 387, 439, 447, 516
9 scenes	50	50 of 100 parent episodes

The model can request saved images from a supplied catalog, subject to workflow selection limits. Structured event updates are validated and persisted by code, including prerequisite-order checks. These supplied records and constraints support the temporal judgments analyzed in Section 5.3.

The action controller applies depth and failed-heading restrictions, can select alignment or recovery actions, and checks completion before accepting STOP. We distinguish model proposals from executed actions using the saved receipts. Offline GT trajectories, endpoints, and benchmark distance-to-goal are excluded from online decisions.

3.3. Model and Execution Settings

The saved request identifies the model as gpt-6-astra, accessed through the OpenAI-compatible chat interface with `reasoning_effort=medium`. The recorded output-token limit is 8,192 per generation. Initial observation provides 12 panoramic headings. Online actions are forward movement of 0.25 m, left/right rotation of 15°, and STOP, under a 100-step episode budget. The configuration records seed 0; this does not establish deterministic API behavior. We analyze one execution of this configuration, without repeated trials, model substitutions, or component ablations. Here, “zero-shot” means no navigation-specific fine-tuning of GPT-6-Astra; the subset was not held out from all workflow development.

3.4. Metrics and Evidence

We report terminal navigation error (NE), success rate (SR), oracle success rate (OSR), success weighted by path length (SPL), and normalized dynamic time warping (nDTW) [12, 13]. These are system-level outcomes. The success radius is 3 m; OSR tests whether the trajectory ever reaches

that neighborhood, whereas SR concerns its endpoint. SPL and nDTW summarize path efficiency and reference-path agreement, respectively. The saved evaluator also counts successful endpoints reached at the budget limit. We use *autonomous completion* to denote a successful endpoint accompanied by a workflow-accepted STOP; this remains a system-level outcome. Evidence comprises final metrics, original images, model outputs, executed-action receipts, and offline GT. Action index t denotes the action from observation frame t to $t + 1$. Section 5 combines analyses of all evaluated episodes with cases that illustrate the observed patterns. Its labels indicate what the saved evidence supports; they do not measure isolated model accuracy.

4. Quantitative Results

4.1. Overall Navigation Performance

Across all 50 evaluated episodes, the GPT-6-Astra-based system achieves an NE of 4.34 m, nDTW of 70.79%, OSR of 62.00%, SR of 52.00%, and SPL of 48.90% (Table 2). These results summarize one recorded execution on the fixed subset in Table 1: 26 episodes end within the success neighborhood. They characterize the complete observation, context-management, decision, and execution workflow described in Section 3.

Table 2 organizes published results along three independent attributes: learning regime, evaluation cohort, and observation setting. Full denotes all 1,839 R2R-CE val-unseen episodes; ON-100 denotes the 100-episode pool used by Open-Nav and subsequent work [19, 20]. Other sampled or unspecified cohorts are identified separately. Pano. and Mono. distinguish panoramic and monocular observations, while F3+P denotes three forward-facing views with on-demand panoramas, as used by Fast-SmartWay and C²Nav [21, 36]. The evaluated system uses frontal RGB-D observations after an initial panorama and is listed under Mono. The evaluated subset uses the corrected R2R_VLNCE_v1-3 release. Differences in episode coverage, data versions, observation access, and action budgets make the table a contextual comparison rather than a matched ranking.

Among zero-shot ON-100 results, Three-Step Nav reports 34.00% SR and 29.12% SPL with panoramic observations [37]. Fast-SmartWay and C²Nav report 27.75%/24.95% and 44.00%/29.00% SR/SPL, respectively, with F3+P observations; the former averages four runs. AgenticNav with GPT-5.5 reports 55.00% SR and 48.41% SPL [15]. The evaluated system’s SPL is numerically close to the latter result, while its SR is lower. The different cohorts and workflows preclude attributing these differences to the language model or observation interface alone.

Supervised results provide a further reference under their own evaluation conditions. On the full split, panoramic

Table 2. **R2R-CE val-unseen navigation results.** Val unseen denotes the cohort; View denotes the observation setting. Our GPT-6-Astra-based system uses 50 episodes from the Open-Nav 100-episode pool [19]. NE is in meters; other metrics are percentages. Cohorts and workflows differ, so these results are not a matched model ranking.

Method	Source	Val unseen	View	NE ↓	nDTW ↑	OSR ↑	SR ↑	SPL ↑
<i>VLN-CE supervised learning</i>								
CMA [11]	CVPR’2022	ON-100	Pano.	6.92	50.77	45.00	37.00	32.17
RecBERT [11]	CVPR’2022	ON-100	Pano.	5.80	54.81	57.00	48.00	43.22
BEVBert [1]	ICCV’2023	ON-100	Pano.	5.13	61.40	64.00	60.00	53.41
ETPNav [2]	TPAMI’2025	ON-100	Pano.	5.16	61.16	63.00	58.00	52.19
ScaleVLN [26]	ICCV’2023	Full	Pano.	4.80	–	–	55.00	51.00
ETPNav [2]	TPAMI’2025	Full	Pano.	4.71	–	65.00	57.00	49.00
BEVBert [1]	ICCV’2023	Full	Pano.	4.57	–	67.00	59.00	50.00
HNR [27]	CVPR’2024	Full	Pano.	4.42	–	67.00	61.00	51.00
Energy [16]	NeurIPS’2024	Full	Pano.	4.69	–	65.00	58.00	50.00
g3D-LF [25]	CVPR’2025	Full	Pano.	4.53	–	68.00	61.00	52.00
NaVid [33]	RSS’2024	Full	Mono.	5.47	–	49.10	37.40	35.90
Uni-NaVid [34]	RSS’2025	Full	Mono.	5.58	–	53.30	47.00	42.70
NaVILA [8]	RSS’2025	Full	Mono.	5.22	–	62.50	54.00	49.00
Aux-Think [23]	NeurIPS’2025	Full	Mono.	5.88	–	54.90	49.70	41.70
Dynam3D [28]	NeurIPS’2025	Full	Mono.	5.34	–	62.10	52.90	45.70
StreamVLN [30]	ICRA’2026	Full	Mono.	4.98	–	64.20	56.90	51.90
D3D-VLP [29]	CVPR’2026	Full	Mono.	4.73	–	67.20	61.30	56.10
<i>Zero-shot VLN-CE</i>								
HSGM [14]	CVPR’2026	Full	Pano.	5.42	–	58.70	47.90	32.80
LXMERT [11]	CVPR’2022	ON-100	Pano.	10.48	18.73	22.00	2.00	1.87
MapGPT-CE-GPT4o [†] [4, 20]	ACL’2024	ON-100	Pano.	8.16	–	21.00	7.00	5.04
DiscussNav-GPT4 [†] [17, 20]	ICRA’2024	ON-100	Pano.	7.77	42.87	15.00	11.00	10.51
NavGPT-CE-GPT4 [†] [9, 38]	AAAI’2024	ON-100	Pano.	8.37	–	26.90	16.30	10.20
Open-Nav-GPT4 [19]	ICRA’2025	ON-100	Pano.	6.70	45.79	23.00	19.00	16.10
Open-Nav-Gemini-2.5-pro [9, 19]	ICRA’2025	ON-100	Pano.	7.28	49.51	30.00	23.00	19.90
Open-Nav-Gemini-3-Flash ^r [19]	ICRA’2025	ON-100	Pano.	6.42	54.82	42.00	32.00	27.65
Three-Step Nav-GPT-5 [37]	AISTATS’2026	ON-100	Pano.	5.87	57.70	39.00	34.00	29.12
LaViRA-GPT-4o [10]	ICRA’2026	ON-100	Pano.	6.43	–	43.30	36.00	28.30
LaViRA-Gemini-2.5-pro [10]	ICRA’2026	ON-100	Pano.	6.54	–	48.70	38.30	28.30
EvoNav-Gemini-2.5-pro [9]	CVPR’2026	ON-100	Pano.	5.04	62.38	51.00	43.00	37.77
SmartWay-GPT-5.5 ^r [20]	IROS’2025	ON-100	Pano.	5.16	58.64	60.00	44.00	35.04
AgenticNav-Gemini-2.5-pro [15]	arXiv’2026	ON-100	Pano.	5.91	48.73	63.00	49.00	33.20
AgenticNav-GPT-5.5 [15]	arXiv’2026	ON-100	Pano.	5.19	63.41	65.00	55.00	48.41
SpatialAnt ^p [35]	arXiv’2026	S-100	Pano.	4.42	69.50	76.00	66.00	54.40
HarnessVLN-GPT-5.5 [7]	arXiv’2026	–	Pano.	4.01	–	72.70	60.80	43.50
Fast-SmartWay-GPT-4o [21]	arXiv’2025	ON-100	F3+P	7.72	51.83	–	27.75	24.95
C ² Nav-GPT-5.5 [36]	arXiv’2026	ON-100	F3+P	7.20	–	54.00	44.00	29.00
CA-Nav [6]	TPAMI’2025	Full	Mono.	7.58	–	48.00	25.30	10.80
AO-Planner [5]	AAAI’2025	Full	Mono.	6.95	–	38.30	25.50	16.60
DreamNav [24]	arXiv’2025	613 traj.	Mono.	7.06	–	41.00	32.80	29.00
GC-VLN [32]	CoRL’2025	Full	Mono.	7.30	–	41.80	33.60	16.30
GPT-6-Astra + our workflow	Report’2026	50 of ON-100	Mono.	4.34	70.79	62.00	52.00	48.90

Evaluation sets. Full: all 1,839 val-unseen episodes. ON-100: the reported Open-Nav 100-episode protocol. S-100: 100 sampled episodes with identity to ON-100 unverified. 50 of ON-100: the evaluated subset (Table 1). A dash denotes an unspecified episode cohort. DreamNav reports 613 trajectories; instruction multiplicity is unspecified. LaViRA scores are means over three runs; both variants use Qwen2.5-VL-32B for visual grounding. Fast-SmartWay scores are means over four runs. Missing metrics remain “–”; bold identifies this report’s system result.

Inputs and settings. Pano.: panoramic observations; Mono.: a monocular egocentric camera; F3+P: three forward-facing views plus on-demand panoramas. Fast-SmartWay also uses an initial panorama. ^pSpatialAnt reconstructs scenes and uses simulator depth for R2R-CE waypoint prediction. [†]Extra annotations/oracle signals in the source comparison. ^rReproduced by AgenticNav’s authors.

Table 3. **Endpoint success and workflow-accepted stopping.** Counts cover all 50 episodes. Success uses the endpoint distance criterion; accepted STOP combines model assessment with workflow checks. Budget termination can also yield a successful endpoint.

Termination	Success	Failure	Total
Workflow-accepted STOP	18	3	21
Budget termination	8	21	29
Total	26	24	50

g3D-LF reports 61.00% SR and 52.00% SPL, while monocular D3D-VLP reports 61.30% and 56.10% [25, 29]. These papers also supply the full-split baseline blocks in Table 2; sampled supervised results remain separate. The numerical proximity of selected scores is insufficient to establish that this system matches a trained navigator across the full benchmark.

Within the evaluated subset, 31 trajectories enter the success neighborhood at least once, but five end outside it. This 10-percentage-point OSR–SR gap shows that reaching the goal region does not ensure a successful endpoint. We next examine whether successful endpoints coincide with an accepted stopping decision.

4.2. Arrival and Stopping Outcomes

Successful endpoints do not always coincide with workflow-accepted stopping (Table 3). The workflow accepts STOP in 21 episodes: 18 succeed and three fail the distance criterion. The remaining 29 episodes terminate at the action budget, including eight successful endpoints and 21 failures. Successful navigation with a workflow-accepted STOP thus occurs in $18/50 = 36.0\%$ of the evaluated episodes; budget-terminated successes contribute another 16.0 percentage points to the reported 52.0% SR. Accepted stopping combines model-based arrival assessment with workflow checks, so this measure is not the model’s standalone completion accuracy.

The unsuccessful accepted STOPs occur in EP7, EP371, and EP377, at NE values of 3.04, 3.37, and 4.07 m. Their distances show varying degrees of endpoint error, including one close to the 3 m threshold. Conversely, EP423 ends only 0.68 m from the goal at budget exhaustion while its last online decision still requests forward movement. These outcomes illustrate the completion gap in the third finding: stopping outside the success region or reaching it without an accepted completion decision. The route and observation evidence in Section 5.2 further examines the latter case.

4.3. Efficiency and Inference Workload

Path efficiency and execution workload capture different aspects of navigation. The system records 4,040 exe-

cuted steps, averaging 80.8 steps and 8.97 m of translational path per episode. These comprise 1,808 forward actions, 2,182 rotations, and 50 terminal STOPs, including budget-generated STOPs. Rotations account for 54.0% of the action count. They can support alignment and evidence acquisition while adding no translational path length, so the reported SPL of 48.90% does not imply low action or inference cost.

The traces contain 9,208 recorded model generations, including recorded format-rejected outputs. This corresponds to 2.28 generations per executed step on average, with a median of 205.5 generations per episode and a range of 39–289. Figure 3 relates these counts to executed steps and separates input-token usage by call role. Progress handoff and arrival review together consume 80.2% of recorded input tokens, showing that maintaining progress and assessing completion dominate this workflow’s input workload.

Available usage fields total 110.41 million input tokens and 3.24 million output tokens; reasoning-token metadata are not added again to the completion total. Recorded generation latencies sum to 22.55 hours, a cumulative call-time measure rather than end-to-end runtime. The logs do not establish an attributable bill or complete cache-discount history. These measurements quantify the workload of the evaluated system without isolating which reviews were necessary. Section 5 examines what the model establishes from the available evidence and whether subsequent actions advance the task.

5. Behavioral and Failure Analysis

Five overlapping perspectives examine the strengths and limitations behind the three findings (Figure 1). Sections 5.1–5.3 relate target and event judgments to observations, supplied history, and executed motion. Sections 5.4–5.5 examine information seeking and judgment revision. Across these analyses, we follow supported judgments through subsequent movement and completion outcomes. Table 4 summarizes observed strengths, unresolved difficulties, and evidence boundaries. Model responses show expressed judgments, external records show retained context, and execution receipts show what physically happened. These perspectives do not define independent capability scores.

5.1. Visual–Language Grounding

Identifying the destination requires more than recognizing its category. We examine whether the saved evidence supports both the stated object or scene category and the qualifiers that identify the intended destination. The analysis covers all 50 completed episodes, with one *terminal destination reference* per original instruction. Repeated observations of the same target count once; intermediate landmarks are outside this inventory. The preliminary offline,




Figure 3. **Inference workload across all 50 episodes.** Left: recorded input tokens by call role. Right: model generations versus executed steps; green denotes endpoint success and red denotes failure. Multiple reviews can precede one action, and tokens include repeated context. Counts describe recorded workload rather than billed cost.

Table 4. **Behavioral evidence within the evaluated workflow.** Five overlapping perspectives connect contextual model judgments with observed execution. Cases and process audits describe this system; they do not isolate intrinsic model capabilities.

Perspective	Observed strengths	Limitations and evidence boundaries
Visual–language grounding	EP218 distinguishes a second-left doorway from a separate forward opening; EP423 identifies the bathtub.	EP469 observes relevant-looking display panes while the route-qualified destination remains unresolved.
Spatial judgments and motion	EP423 follows the two-hallway route; EP218’s translated viewpoint exposes an occluded doorway.	EP116 does not complete a visible doorway crossing; executed movement is also constrained by the controller.
Temporal reasoning with history	Given retained records and observations, EP423 relates later confirmation to an earlier turn; EP116 distinguishes turning from crossing.	These judgments use supplied history and order checks; they do not establish unaided memory.
Action adjustment	Across the cohort, 36 of 66 audited problem-handling processes show local relief, covering 24 episodes.	29 processes remain unresolved; local relief does not isolate a model contribution or guarantee navigation success.
Evidence-guided verification	48 of 75 audited verification processes clarify an uncertain judgment with supporting evidence, covering 33 episodes.	22 processes retain uncertainty and five are unassessable; the workflow initiates reviews.

assistant-assisted audit combines saved target judgments, execution receipts, GT/executed paths, and 319 selected original RGB frames across the nine scenes. The judgments are conditioned on workflow-selected views and context. Model confidence and internal “supported” flags are claims to inspect, not annotation labels; GT is used only offline.

We distinguish (1) *category recognition*, whether a visually identified candidate has the stated object or scene category; (2) *relational grounding*, whether observations support its instruction-qualified spatial or route relation; and (3) *target-instance correspondence*, whether the selected physical reference is consistent with both the visual context and the GT destination route. Each dimension is labeled supported (S), reference-inconsistent (C), or

unassessable (U). U includes targets not established in the recorded views, unresolved alternatives, and insufficient independent evidence. A target not reached or not identified is therefore not automatically a recognition error. Composite references require their relevant components: seeing a bathroom, for example, does not establish the separate toilet room named in the instruction.

The intended target is established less often than its category. Table 5 reports support in 33/50 episodes for category recognition, 23/50 for relations, and 18/50 for target-instance correspondence. Fifteen episodes have category support without independently supported instance correspondence. This gap includes ambiguity and limited visibility as well as an observed candidate/reference-route mis-

Table 5. **Terminal-reference evidence across all 50 episodes.** Preliminary offline assistant-assisted audit, one reference per episode. S: supported; C: reference-inconsistent; U: unassessable. The last column uses all 50 episodes, not only assessable cases.

Dimension	S	C	U	S / 50
<i>Category recognition</i>	33	0	17	66.0%
<i>Relational grounding</i>	23	0	27	46.0%
<i>Target-instance correspondence</i>	18	1	31	36.0%

Assessable counts (S+C) are 33, 23, and 19. U includes targets not established in the saved views and unresolved reference identity. C denotes a candidate/GT-route inconsistency, not an isolated perceptual cause. These are evidence-coverage counts, not validated accuracy scores.

match; it is not a count of 15 semantic errors. The 31 unassessable instance judgments remain in the denominator. These percentages measure evidence coverage within this cohort, not validated capability accuracy. Different visibility and assessability also prevent interpreting their difference as a calibrated ranking of the three abilities. No independent human annotation or inter-rater study has been performed.

Qualitative grounding evidence. Figure 4 links instruction phrases, visual evidence, and trajectories across three settings. In panel (a), EP218 distinguishes the “second on the left” from neighboring openings; EP348 separates a balcony door from adjacent windows, and EP387 establishes passage between the stove and island and beyond the island. Visual and route evidence jointly support all three grounding dimensions. In panel (b), EP469 recognizes the glass information panes, but the selected display lies on a different branch from the GT destination route while the doorway qualifier remains unresolved. This is a reference-route inconsistency, with its cause undetermined. EP226 and EP247 likewise recognize the relevant category but do not resolve the required entry order or hallway-turn relation; their instance labels are U, whereas EP469 alone is C. In panel (c), EP116 stalls before observing the terminal tub and sink, EP40 remains in the kitchen-side area without observing the bedroom bed, and EP232 leaves the sculpture unresolved among candidate objects. All three dimensions are U for these cases. Supported cases illustrate the first finding; unresolved cases show why category recognition alone is insufficient to identify the instructed target. Missing evidence remains distinct from recognition error.

5.2. Spatial Judgments and Executed Motion

Seeing a passage, crossing it, and following the instructed route are separate steps. Three perspectives organize the cases: (1) *position and heading*, relative to an entrance or travel lane; (2) *local passage*, relating visible openings to depth, obstacles, and movement constraints; and (3) *spatial connectivity*, distinguishing an onward corridor from

a side room or another branch. These overlap with the instruction-qualified relations in Section 5.1. We examine model judgments, controller-selected or constrained actions, and recorded motion together to characterize progress and stalled crossings; they do not isolate spatial-perception accuracy.

From facing a passage to following its connection. Figure 5 makes these distinctions visible. In panel (a), EP13 turns 90° from a wall toward the kitchen aisle, and EP166 turns 45° toward the window-side passage; both pairs share a position, isolating the change in heading. EP166’s initial alignment is executed by the panorama-bootstrap controller, so the examples characterize the complete workflow. In panel (b), EP387 advances 2.75 m between the stove and island without contact. EP116 instead remains at the same doorway position: Step 35 has 0.715 m center depth but a failed-heading restriction, whereas Step 60 has only 0.205 m center depth and a depth restriction. The visible bedroom does not establish a completed crossing, and these records do not show that bypassing either restriction would be safe. In panel (c), EP423’s views expose a continuing corridor beside a separate room, consistent with its reference-route alignment; its episode nevertheless ends at the budget limit. EP469 moves without contact along another gallery branch, leaving the instructed route unresolved. Together, the cases connect viewpoint geometry, admissible motion, and spatial connectivity without treating a recognizable destination or collision-free movement as sufficient route understanding.

5.3. Temporal Reasoning with Supplied History

GPT-6-Astra’s recorded reviews include distinctions between completed and unfinished events under the supplied history. We examine these judgments through (1) *event timing*, separating physical occurrence from later confirmation; (2) *order constraints*, distinguishing prerequisites such as entering a room before the next turn; and (3) *event continuity*, explicitly referring to a past occurrence while distinguishing a later visit or a different instruction match. The workflow retains event records, selects historical views and receipts, and applies prerequisite-order checks. The analysis therefore concerns judgments made with this support, not autonomous retention of the navigation history.

Evidence across all evaluated episodes. We conduct a preliminary offline, assistant-assisted audit of all 50 episodes. The original instructions yield 172 motion clauses and 122 adjacent clause pairs, including unreached clauses; terminal waiting and stopping are treated separately. Saved model responses are checked against original views and execution receipts. Timing checks whether a stated occurrence and interval are compatible with the recorded motion, without validating an exact earliest boundary or the intended route. Continuity uses 132 clauses whose first oc-

(a) Supported correspondence

Category, relation and instance supported

EP218 Instruction **Goal: D2**

room ... **second on the left**
Stop in the **doorway**



Category S · Relation S · Instance S

EP348 Instruction **Goal: B1**


first door to the balcony
on your left



Category S · Relation S · Instance S

EP387 Instruction **Landmark: I1**

between the **stove** and the **island**
past the **island**




Category S · Relation S · Instance S

(b) Uncertainty / mismatch

Category supported; instance C or U

EP469 Instruction **Candidate: P1**


Walk **past** the **doorway**
wait by the **glass info panes**



Category S · Relation U · Instance C

EP226 Instruction **Candidates: A / B**


third entry way on your left
enter the **bedroom**



Category S · Relation U · Instance U

EP247 Instruction **Candidate: C1**

at the end of hallway
turn slight left ... Stop at **couch**




Category S · Relation U · Instance U

(c) Target not established

Terminal-reference evidence unassessable

EP116 Instruction **Target: ?**


walk into bathroom
near the **tub** and **sink**



Category U · Relation U · Instance U

EP40 Instruction **Target: ?**


Turn left ... into the **bedroom**
corner of the **bed**



Category U · Relation U · Instance U

EP232 Instruction **Candidates: A? / B?**

room at the end of the hall
wait at the **sculpture**



Category U · Relation U · Instance U

— Category region / word — Relation cue / word ID Candidate instance -- Uncertain / contextual region
— Executed — GT route ● Start ★ GT goal ■ Endpoint ○ Shown step


Figure 4. **Visual-language grounding in navigation.** Panels show supported correspondence (a), uncertain or inconsistent reference (b), and targets not established (c). Blue, ochre, and purple link category words/regions, relational cues, and candidate IDs, respectively. Step t denotes the observation after t actions. S/C/U follow Table 5.

(a) Where am I facing?


Position and heading

EP13 Instruction

continue towards the kitchen

**EP166** Instruction


Go forward by the windows

**(b) Can I move through?**


Local passage

EP387 Instruction

between the stove and the island

**EP116** Instruction


walk out of closet into a bed room

**(c) Where does it lead?**


Spatial connectivity

EP423 Instruction

walk across this hallway

**EP469** Instruction

Walk past the doorway



— Earlier-view heading — Later-view heading — Navigable surface — Depth sampling window
— Executed — GT route ● Start ★ GT goal ■ Endpoint

Figure 5. **Spatial judgments and executed motion.** Columns show position and heading (a), local passage (b), and spatial connectivity (c), including controller-mediated movement. Paired RGB views match the labeled poses. Image arrows highlight passage cues; map arrows show heading. Dashed boxes mark the depth sampling window; values are its recorded 10th-percentile depth. Navmesh crops and GT routes are offline references. Local maps include scale bars; full trajectories use independent scales.

Table 6. **Temporal judgments under supplied context.** The audit covers all 50 episodes. N counts motion clauses, adjacent order pairs, or eligible continuity checks. S: supported; C: contradicted; U: unassessable.

Dimension	N	S	C	U	S/N
<i>Event timing</i>	172	146	0	26	84.9%
<i>Order constraints</i>	122	100	0	22	82.0%
<i>Event continuity</i>	132	61	0	71	46.2%

Timing covers 50 episodes; order covers 44 with multiple motion clauses; continuity covers 47 with a later review at least 10 actions after the first occurrence claim. Repeated answers count once per unit. S/N measures evidence coverage, not accuracy.

currence claim has a progress-review opportunity at least ten actions later. It requires an explicit reference to the identifiable past event; an unchanged controller record alone is insufficient. Repeated answers about one unit count once, and the model’s own support flags are not correctness labels.

Table 6 summarizes supported evidence for the instruction clauses and pairs defined above. Averaging S/N within each eligible episode gives coverage of 85.6%, 86.2%, and 52.8% for timing, order, and continuity, respectively, with each eligible episode receiving equal weight.

Distribution across episodes. Figure 6 complements the unit-level census by grouping episodes according to their applicable checks. All S means every applicable check has supported evidence; Mixed S/U combines supported and unassessable checks; All U means every applicable check is unassessable; and N/A denotes no applicable check.

Timing and order each have 31 All S episodes, compared with 14 for continuity. A further 17, 12, and 22 episodes, respectively, contain mixed evidence. Thus, the supported judgments extend across many episodes, while complete continuity coverage is less common. The six episodes without an applicable order pair and the three without an eligible continuity check remain visible as N/A, so every column retains the full cohort of 50.

In EP423, later views confirm an earlier turn’s instruction match, without treating the confirmation as a new turn. In EP116, reviews distinguish the completed turn from the pending crossing, consistent with the lack of translation. These cases connect temporal judgments to the first and third findings: clarifying past actions and recognizing unfinished ones without completing them.

Unassessable checks include ambiguous boundaries, unexercised instruction events, and missing explicit later judgments. Continuity also has stricter eligibility and review requirements, so the group sizes do not rank temporal abilities. No contradiction was established in this preliminary audit; All U records insufficient evidence, not an observed temporal error. Supported responses show compatibility with recorded events under the supplied context. Without




Figure 6. **Temporal judgments with supplied history.** All 50 episodes are grouped by applicable checks: all supported (All S), mixed supported/unassessable (Mixed S/U), all unassessable (All U), or none (N/A). Counts measure evidence coverage under external context management, not unaided memory.

varying history selection or record retention, the counts do not separate model reasoning from the assistance provided by that context.

5.4. Action Adjustment and Local Outcomes

Some viewpoint changes reveal previously hidden passages. We examine all 50 episodes using saved model proposals, RGB observations, and matching execution records. Proposals are interpreted alongside the workflow’s movement constraints and selected actions. The preliminary offline, assistant-assisted audit identifies 66 problem-handling processes in 42 episodes; eight episodes contain no eligible process. Repeated attempts addressing the same problem are merged, while routine instruction following is excluded. Each process is grouped by its initial audited adjustment, and local relief requires evidence addressing the original difficulty, such as exposing an obscured passage or clearing an obstruction. These are observed system processes, not isolated tests of the model’s planning contribution.

The ring in Figure 7(a) summarizes the strategy mix: changes of viewing angle account for 33 processes (50.0%), and adjustments to blocked approaches account for 27 (40.9%). The remaining processes inspect another candidate (four, 6.1%) or return along the earlier route (two, 3.0%). The bars below show the corresponding local outcomes, including 15 relieved processes in the first group and 18 in the second. Across all groups, 36 processes show local relief, 29 remain unresolved at the observed end, and one is unassessable. These counts do not rank strategy effectiveness: a process may include several later adjustments, the groups encounter different problems, and the latter two groups are small.

Figure 7(b) illustrates how an adjustment changes the available evidence. In EP244, stationary left/right inspections leave the landing’s left continuation hidden by a nearby door leaf. At Step 23, GPT-6-Astra proposes a short forward movement to see past it. The recorded 0.254 m translation exposes the left-side floor at Step 24; three left turns then align the view with the passage at Step 27. The useful outcome is the newly visible passage that answers the model’s inspection question. This example illustrates the visual information seeking in the second finding; the process counts show both local relief and unresolved difficulties. Local relief remains distinct from final navigation success, and the recorded workflow does not establish a causal gain over continuing the previous strategy.

5.5. Evidence-Guided Verification

Some reviews resolve uncertainty through additional visual evidence. The workflow initiates progress and arrival reviews and supplies earlier records. GPT-6-Astra’s recorded responses include questions, requests for saved images, and revised or retained judgments. Our preliminary assistant-assisted audit follows these verification processes across all 50 episodes, linking an earlier judgment, an attempt to obtain discriminating evidence, and a later judgment about the same reference or interpretation. Repeated reviews of an unresolved question form one process. Conclusions are checked against saved observations and execution records; the model’s own support flag does not determine the audit label.

Figure 8(a) summarizes 75 verification processes in 48 episodes. Of these, 48 (64.0%) resolve an initially uncertain judgment into an evidence-supported conclusion, including both identifying and justifiably rejecting a candidate. Another 22 (29.3%) retain uncertainty, while five (6.7%) have an unassessable endpoint. The episode bar counts each episode once: 33 contain at least one supported clarification, 15 contain eligible processes but no supported clarification, and two contain no eligible process. A local clarification does not imply that every question in that episode is resolved.

Figure 8(b) illustrates a supported rejection in EP308. At Step 40, GPT-6-Astra questions whether the partly visible opening leads to the instructed hallway. Further views expose shelving and a back wall; at Step 66, it explicitly rejects the opening as a hallway candidate. This resolves one competing interpretation while leaving the overall hallway route uncertain.

These processes show strengths in revising local judgments, while sometimes leaving the correct route unresolved. All independently assessable pairs begin with uncertainty, so this inventory provides no denominator for an error-correction rate. In the 22 unresolved processes, uncertainty persists through the observed reviews.

These results describe judgment revision within workflow-scheduled calls; they do not establish spontaneous initiation of reflection or a causal benefit over a system without these reviews.

6. Discussion and Limitations


Later views support a match between an earlier action and the instruction. In EP423, later views allow GPT-6-Astra to confirm the earlier turn’s role in the instructed route while retaining when it occurred. Across the evaluated episodes, all applicable timing and order checks have supported evidence in 31 episodes each, although timing checks alone do not establish a route match. The time of an action and the time when its instruction match becomes clear need not coincide. Assessing this distinction requires inspecting both the new observation and the supplied history, including the workflow’s event records and order checks.

Additional observations support local clarification in some cases. In EP244, a proposed forward movement exposes a hidden passage; in EP308, further views support rejecting an uncertain hallway candidate. Across all evaluated episodes, 36 of 66 adjustment processes show local relief and 48 of 75 verification processes produce a supported clarification. These findings suggest checking what an observation reveals, then whether that information changes navigation. Exposing a passage can be useful before advancing along it, whereas repeated reviews may leave the same question unresolved. The two inventories use different units and eligibility rules; their proportions do not measure a common success rate or establish causal navigation gains.

The gap between understanding and completion spans movement and stopping. In EP116, GPT-6-Astra distinguishes a completed turn from an unfinished crossing while the system continues rotating. In EP308, rejecting one opening leaves the route unresolved. At termination, the 52.0% endpoint success rate includes 16.0 percentage points from budget termination, while three accepted STOPS fail the distance criterion. These distinct outcomes share an evaluation requirement: follow the judgment through the selected action, actual motion, and stopping decision. The recorded sequence can locate where progress stops; identifying which model or controller choice caused the failure requires controlled comparisons.

The results depend on both the model and the workflow. External code selects context, retains records, schedules reviews, and constrains actions. Requests receive selected earlier outputs without replaying full conversations or reasoning-state items. Trajectories measure system performance; responses show judgments under supplied context. Neither isolates unaided memory, spontaneous reflection, or intrinsic navigation ability. The data do not estab-

(a) Strategy mix and local outcomes



(b) EP244: moving past a visual obstruction

Instruction: "Go through the door and turn left."





Figure 7. **Action adjustments and local outcomes.** (a) The ring shows initial-strategy shares among 66 audited system processes; bars show local outcomes, with totals n at the right. Hatching denotes unassessable outcomes. (b) EP244's executed viewpoint change reveals a passage hidden during stationary inspection. Gold outlines the door leaf; blue indicates the revealed passage. Plan text paraphrases the model proposal. Step t is the observation after t actions.

(a) Verification across the cohort



(b) EP308: testing a candidate passage

Instruction: "Leave the bedroom and go into the hallway."




Figure 8. **Evidence-guided verification in scheduled reviews.** (a) Outcomes of 75 processes and mutually exclusive coverage groups across all 50 episodes. Supported clarification includes justified candidate rejection. (b) EP308 resolves one opening as storage rather than the instructed hallway. Ochre outlines track the same candidate; judgment summaries paraphrase saved model reviews. Step t denotes the observation after t actions.

lish whether these workflow choices assist or restrict performance; retaining more context or removing a movement restriction is not demonstrated to help.

Supported evidence is not a model accuracy score. The audits cover all 50 episodes but use different units: terminal references for grounding, instruction events for

temporal judgments, and eligible adjustment or verification processes. Spatial analysis remains case-based. These measurements cannot be combined into a single capability score. Unassessable labels retain missing or ambiguous evidence rather than treating it as an error. The retrospective, assistant-assisted annotations have not undergone blinded human review or inter-rater validation, and model explanations are reported outputs rather than access to internal reasoning. Offline GT helps check route correspondence; it does not establish the outcome of an unexecuted alternative.

Broader claims require fresh episodes and matched comparisons. The evaluated subset is a fixed 50-episode prefix spanning nine scenes, with one recorded run per episode and exposure during workflow development. It does not estimate run-to-run variability or full-split performance. Published results with different cohorts, data versions, observation access, budgets, and stopping rules provide context without a controlled superiority comparison. Future evaluation should freeze these conditions, use fresh episodes and repeated runs, and independently validate the annotations. Holding the workflow fixed while substituting models would compare their contributions within that workflow; holding the model fixed while varying context retention would test the effect of context management. Following clarified judgments through subsequent actions would further test whether local resolution persists into route progress and appropriate stopping.

7. Conclusion

The evaluated GPT-6-Astra-based navigation system achieves 52.0% SR, 48.9% SPL, and 70.8% nDTW with external context management and action control. Its recorded responses show strengths in linking landmarks and earlier actions to instructions, seeking visual information, and revising uncertain judgments. The observed gap between task understanding and autonomous completion includes continued rotation after an unfinished crossing is recognized. At termination, eighteen episodes succeed with a workflow-accepted STOP, eight succeed at the budget limit, and three accepted STOPS fail the distance criterion. These findings concern the complete system and make the next challenge concrete: translating supported judgments into the required movement, sustained route progress, and appropriate stopping.

References

- [1] Dong An, Yuankai Qi, Yangguang Li, Yan Huang, Liang Wang, Tieniu Tan, and Jing Shao. Bevbort: Multimodal map pre-training for language-guided navigation. In *ICCV*, pages 2737–2748, 2023.
- [2] Dong An, Hanqing Wang, Wenguan Wang, Zun Wang, Yan Huang, Keji He, and Liang Wang. ETPNav: Evolving topological planning for vision-language navigation in continuous environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(7):5130–5145, 2025.
- [3] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian D. Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, pages 3674–3683, 2018.
- [4] Jiaqi Chen, Bingqian Lin, Ran Xu, Zhenhua Chai, Xiaodan Liang, and Kwan-Yee 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.
- [5] Jiaqi Chen, Bingqian Lin, Xinmin Liu, Lin Ma, Xiaodan Liang, and Kwan-Yee K. Wong. Affordances-Oriented Planning Using Foundation Models for Continuous Vision-Language Navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 23568–23576, 2025.
- [6] Kehan Chen, Dong An, Yan Huang, Rongtao Xu, Yifei Su, Yonggen Ling, Ian Reid, and Liang Wang. Constraint-Aware Zero-Shot Vision-Language Navigation in Continuous Environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(11):10441–10456, 2025.
- [7] Yang Chen, Lirong Che, Zhenyu Huang, Wenbo Fu, Chuang Wang, Xu Cao, Daqi Liu, Yuzhe Yang, Jian Su, and Lan-Zhe Guo. HarnessVLN: Unifying Training-Free Embodied Navigation through an Agent Harness. *arXiv preprint arXiv:2609.15195*, 2026.
- [8] An-Chieh Cheng, Yandong Ji, Zhaojing Yang, Zaitian Gongye, Xueyan Zou, Jan Kautz, Erdem Biyik, Hongxu Yin, Sifei Liu, and Xiaolong Wang. NaVILA: Legged Robot Vision-Language-Action Model for Navigation. In *Robotics: Science and Systems*, 2025.
- [9] Guangzhao Dai, Shuo Wang, Zihan Wang, Guo-Sen Xie, Yang Yang, Jinshan Pan, Qianru Sun, and Xiangbo Shu. History to future: Evolving agent with experience and thought for zero-shot vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15177–15187, 2026.
- [10] Hongyu Ding, Ziming Xu, Yudong Fang, You Wu, Zixuan Chen, Jieqi Shi, Jing Huo, Yifan Zhang, and Yang Gao. LaViRA: Language-Vision-Robot Actions Translation for Zero-Shot Vision Language Navigation in Continuous Environments. In *IEEE International Conference on Robotics and Automation*, 2026. ICRA 2026 acceptance stated in the authors’ arXiv record; results from version 2.
- [11] Yicong Hong, Zun Wang, Qi Wu, and Stephen Gould. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *CVPR*, pages 15418–15428, 2022.
- [12] Gabriel Ilharco, Vihan Jain, Alexander Ku, Eugene Ie, and Jason Baldridge. General evaluation for instruction conditioned navigation using dynamic time warping. In *NeurIPS Workshop on Visually Grounded Interaction and Language (ViGIL)*, 2019.
- [13] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision-and-

- language navigation in continuous environments. In *ECCV*, pages 104–120, 2020.
- [14] Kailing Li, Tianwen Qian, Lijin Yang, Yuqian Fu, Jingyu Gong, Xiaoling Wang, and Liang He. Bridging the 2D-3D Gap: A Hierarchical Semantic-Geometric Map for Vision Language Navigation. In *CVPR*, 2026.
 - [15] Yijian Li, Changze Li, Hantian Shi, Jiaying Luo, Jiyuan Cai, Ming Yang, and Tong Qin. AgenticNav: Zero-shot vision-and-language navigation as a tool-calling harness. *arXiv preprint arXiv:2606.10577*, 2026.
 - [16] Rui Liu, Wenguan Wang, and Yi Yang. Vision-language navigation with energy-based policy. *Advances in Neural Information Processing Systems*, 37:108208–108230, 2024.
 - [17] Yuxing Long, Xiaoqi Li, Wenzhe Cai, and Hao Dong. Discuss before moving: Visual language navigation via multi-expert discussions. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
 - [18] Yanyuan Qiao, Haodong Hong, Wenqi Lyu, Dong An, Siqi Zhang, Yutong Xie, Xinyu Wang, and Qi Wu. Navbench: Probing multimodal large language models for embodied navigation. *arXiv preprint arXiv:2506.01031*, 2025.
 - [19] Yanyuan Qiao, Wenqi Lyu, Hui Wang, Zixu Wang, Zerui Li, Yuan Zhang, Minghui Tan, and Qi Wu. Open-nav: Exploring zero-shot vision-and-language navigation in continuous environment with open-source llms. In *ICRA*, pages 6710–6717, 2025.
 - [20] Xiangyu Shi, Zerui Li, Wenqi Lyu, Jiatong Xia, Feras Dayoub, Yanyuan Qiao, and Qi Wu. Smartway: Enhanced way-point prediction and backtracking for zero-shot vision-and-language navigation. In *IROS*, pages 16923–16930, 2025.
 - [21] Xiangyu Shi, Zerui Li, Yanyuan Qiao, and Qi Wu. Fast-smartway: Panoramic-free end-to-end zero-shot vision-and-language navigation. *arXiv preprint arXiv:2511.00933*, 2025.
 - [22] VLN-CE Project. VLN-CE data: R2R_VLNCE_v1-3. <https://jacobkrantz.github.io/vlnce/data>, 2022. Dataset release notes; accessed September 15, 2026.
 - [23] Shuo Wang, Yongcai Wang, Wanting Li, Xudong Cai, Yucheng Wang, Maiyue Chen, Kaihui Wang, Zhizhong Su, Deying Li, and Zhaoxin Fan. Aux-Think: Exploring Reasoning Strategies for Data-Efficient Vision-Language Navigation. In *Advances in Neural Information Processing Systems*, 2025.
 - [24] Yunheng Wang, Yuetong Fang, Taowen Wang, Yixiao Feng, Yawen Tan, Shuning Zhang, Peiran Liu, Yiding Ji, and Renjing Xu. Dreamnav: A trajectory-based imaginative framework for zero-shot vision-and-language navigation. *arXiv preprint arXiv:2509.11197*, 2025.
 - [25] Zihan Wang and Gim Hee Lee. g3D-LF: Generalizable 3D-Language Feature Fields for Embodied Tasks. In *CVPR*, pages 14191–14202, 2025.
 - [26] Zun Wang, Jialu Li, Yicong Hong, Yi Wang, Qi Wu, Mohit Bansal, Stephen Gould, Hao Tan, and Yu Qiao. Scaling data generation in vision-and-language navigation. In *ICCV*, pages 12009–12020, 2023.
 - [27] Zihan Wang, Xiangyang Li, Jiahao Yang, Yeqi Liu, Junjie Hu, Ming Jiang, and Shuqiang 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.
 - [28] Zihan Wang, Seungjun Lee, and Gim Hee Lee. Dynam3D: Dynamic Layered 3D Tokens Empower VLM for Vision-and-Language Navigation. In *Advances in Neural Information Processing Systems*, 2025.
 - [29] Zihan Wang, Seungjun Lee, Guangzhao Dai, and Gim Hee Lee. D3D-VLP: Dynamic 3D Vision-Language-Planning Model for Embodied Grounding and Navigation. In *CVPR*, pages 32463–32474, 2026.
 - [30] Meng Wei, Chenyang Wan, Xiqian Yu, Tai Wang, Yuqiang Yang, Xiaohan Mao, Chenming Zhu, Wenzhe Cai, Hanqing Wang, Yilun Chen, Xihui Liu, and Jiangmiao Pang. StreamVLN: Streaming Vision-and-Language Navigation via SlowFast Context Modeling. In *IEEE International Conference on Robotics and Automation*, 2026. ICRA 2026 acceptance stated in the authors’ arXiv record.
 - [31] Wei Xue, Mingcheng Li, Xuecheng Wu, Jingqun Tang, Dingkan Yang, and Lihua Zhang. ProFocus: Proactive perception and focused reasoning in vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18129–18139, 2026.
 - [32] Hang Yin, Haoyu Wei, Xiuwei Xu, Wenxuan Guo, Jie Zhou, and Jiwen Lu. GC-VLN: Instruction as Graph Constraints for Training-free Vision-and-Language Navigation. In *Proceedings of The 9th Conference on Robot Learning*, pages 1809–1824, 2025.
 - [33] Jiazhao Zhang, Kunyu Wang, Rongtao Xu, Gengze Zhou, Yicong Hong, Xiaomeng Fang, Qi Wu, Zhizheng Zhang, and He Wang. NaVid: Video-based VLM Plans the Next Step for Vision-and-Language Navigation. In *Robotics: Science and Systems*, 2024.
 - [34] Jiazhao Zhang, Kunyu Wang, Shaoan Wang, Minghan Li, Haoran Liu, Songlin Wei, Zhongyuan Wang, Zhizheng Zhang, and He Wang. Uni-NaVid: A Video-based Vision-Language-Action Model for Unifying Embodied Navigation Tasks. In *Robotics: Science and Systems*, 2025.
 - [35] Jiwen Zhang, Xiangyu Shi, Siyuan Wang, Zerui Li, Zhongyu Wei, and Qi Wu. SpatialAnt: Autonomous Zero-Shot Robot Navigation via Active Scene Reconstruction and Visual Anticipation. *arXiv preprint arXiv:2603.26837*, 2026.
 - [36] Runtian Zheng, Congpeng Zhang, and Ying Liu. C²Nav: Compare Before You Commit for Zero-Shot Vision-and-Language Navigation. *arXiv preprint arXiv:2609.15142*, 2026.
 - [37] Wanrong Zheng, Yunhao Ge, and Laurent Itti. Three-Step Nav: A Hierarchical Global–Local Planner for Zero-Shot Vision-and-Language Navigation. In *Proceedings of The 29th International Conference on Artificial Intelligence and Statistics*, pages 4645–4653. PMLR, 2026.
 - [38] Gengze Zhou, Yicong Hong, and Qi Wu. NavGPT: Explicit reasoning in vision-and-language navigation with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7641–7649, 2024.