

FLAME: Learning to Navigate with Multimodal LLM in Urban Environments

Yunzhe Xu¹, Yiyuan Pan², Zhe Liu^{1*}, Hesheng Wang²

¹MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China

²Department of Automation, Shanghai Jiao Tong University, China
{xyz9911, pyy030406, liuzhesjtu, wanghesheng}@sjtu.edu.cn

Abstract

Large Language Models (LLMs) have demonstrated potential in Vision-and-Language Navigation (VLN) tasks, yet current applications face challenges. While LLMs excel in general conversation scenarios, they struggle with specialized navigation tasks, yielding suboptimal performance compared to specialized VLN models. We introduce FLAME (FLAMingo-Architected Embodied Agent), a novel Multimodal LLM-based agent and architecture designed for urban VLN tasks that efficiently handles multiple observations. Our approach implements a three-phase tuning technique for effective adaptation to navigation tasks, including single perception tuning for street view description, multiple perception tuning for route summarization, and end-to-end training on VLN datasets. The augmented datasets are synthesized automatically. Experimental results demonstrate FLAME’s superiority over existing methods, surpassing state-of-the-art methods by a 7.3% increase in task completion on Touchdown dataset. This work showcases the potential of Multimodal LLMs (MLLMs) in complex navigation tasks, representing an advancement towards applications of MLLMs in the field of embodied intelligence.

Code — <https://github.com/xyz9911/FLAME>

1 Introduction

Large Language Models (LLMs) (Achiam et al. 2023; Touvron et al. 2023) have revolutionized the field of embodied intelligence. Vision-and-Language Navigation (VLN) (Anderson et al. 2018), a fundamental task in embodied AI, challenges agents to navigate to a goal following human instructions in indoor or outdoor environments. This task demands sophisticated abilities in instruction comprehension, environmental understanding, and decision-making, which can be effectively managed by LLMs. Recent approaches have integrated LLMs into VLN methods, either by translating visual data into language (Zhou, Hong, and Wu 2024; Qiao et al. 2023; Chen et al. 2024) or by employing Multimodal LLMs (MLLMs) (Zhang et al. 2024; Zhou et al. 2024) for environmental perception.

However, the incorporation of general-purpose LLMs in VLN still faces critical challenges, primarily stemming from

*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

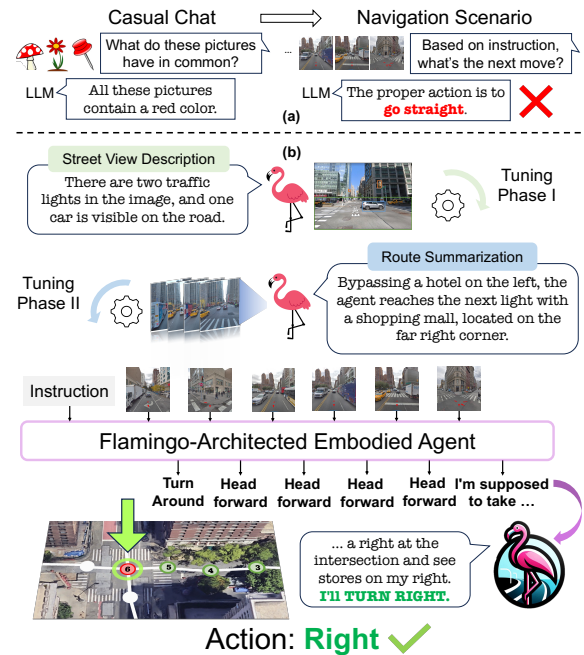


Figure 1: LLM-based agents excel in conversation but often falter in specialized navigation tasks. Our agent, powered solely by a Multimodal LLM, demonstrates proficiency in navigation skills, efficiently adapting to navigation-specific scenarios through targeted finetuning phases.

their inherent limitations in navigation-specific scenarios, as Figure 1(a) depicts. For text-only LLMs, translating visual data into language using visual foundation models can lead to information loss (Zhou, Hong, and Wu 2024), resulting in a performance gap compared to VLN-specialized models. For Multimodal LLMs, while partially addressing the limitations of text-only LLMs, they often struggle to interact within navigation-specific scenarios. Recent attempts to incorporate MLLMs for navigation may have constrained their capabilities by either using them as auxiliary components to traditional VLN models (Zhou et al. 2024) or by processing observations through video tokens (Zhang et al. 2024), which require multiple forward passes per trajectory and significant computational resources in training. These

approaches may not efficiently adapt MLLMs to navigation-specific scenarios, which limits their ability to leverage inherent capabilities for handling interleaved textual and visual inputs and affects overall performance.

Moreover, the application of MLLM in urban VLN remains unexplored, despite its importance alongside indoor VLN. Outdoor navigation presents unique challenges for introducing MLLM, including longer trajectory lengths (up to 55 iterations) and increased difficulty (40% lower success rate compared to indoor navigation tasks).

To address these challenges, we introduce FLAME (FLAMingo-Architected Embodied Agent), the first MLLM-based agent designed for urban VLN tasks, as shown in Figure 1(b). Based on Flamingo (Alayrac et al. 2022), FLAME operates autoregressively and efficiently handles multiple perceptions without increasing context length, ensuring efficiency in end-to-end training and inference. We propose a three-phase tuning technique to adapt Flamingo model (Alayrac et al. 2022) to navigation tasks using augmented data: 1) Single perception tuning: Learning to describe street views. 2) Multiple perception tuning: Learning to summarize agent navigation. 3) End-to-End training and evaluation on VLN datasets. To support the first two tuning phases, we utilize GPT-4 (Achiam et al. 2023) to synthesize captions and route summaries for the Touchdown environment (Chen et al. 2019). Additionally, we synthesize navigation rationales for urban VLN datasets (Chen et al. 2019; Schumann and Riezler 2021) to validate FLAME’s reasoning capability (Wei et al. 2022).

Our agent achieves remarkable computational efficiency, completing the entire training process in just 14 hours on a single A100 GPU. Experimental results demonstrate FLAME’s superiority over existing methods on two urban VLN datasets: Touchdown (Chen et al. 2019) and Map2seq (Schumann and Riezler 2021). Our approach significantly outperforms current state-of-the-art (SOTA) methods by 7.3% Task Completion (TC) in Touchdown and 3.74% TC in Map2seq. Our work not only benefits the field of VLN but also showcases the potential of MLLMs in navigation within complex environments and tasks.

In summary, our contributions are threefold:

- We introduce FLAME, to our knowledge, the first agent based on Multimodal LLM (MLLM) for urban Vision-and-Language Navigation (VLN) tasks.
- We propose a tailored three-phase tuning technique for adapting Flamingo into navigation scenarios using synthetic data, fully unleashing MLLM’s power.
- Experiments show the superiority of FLAME over current SOTAs. FLAME’s performance proves that MLLMs can significantly outperform specialized models, opening new avenues for research in embodied AI.

2 Related Works

Vision-and-Language Navigation

Vision-and-Language Navigation (VLN) (Anderson et al. 2018) encompasses indoor (Qi et al. 2020; Ku et al. 2020) and outdoor scenarios (Chen et al. 2019; Schumann and

Riezler 2021), with most advancements focusing on indoor environments. Traditional VLN agents often lack advanced decision-making skills, prompting the integration of Large Language Models (LLMs) (Lin et al. 2024; Schumann et al. 2024), leveraging the reasoning (Chen et al. 2024) and dialogue capabilities (Qiao et al. 2023; Long et al. 2024) of LLMs. These approaches either convert visual data to text (Zhou, Hong, and Wu 2024) or employ Multimodal LLMs (MLLMs) with intensive training (Zhang et al. 2024; Zheng et al. 2024). Given the relative underexploration of outdoor VLN, our work addresses this gap by introducing an effectively adapted MLLM-based agent for urban VLN tasks.

Multimodal Large Language Models

The emergence of Multimodal Large Language Models (MLLMs) (Liu et al. 2024; Alayrac et al. 2022; Awadalla et al. 2023) has expanded their application to various tasks, including captioning (Li et al. 2023b) and general instruction following (Dai et al. 2023). These models demonstrate multimodal reasoning skills (Lu et al. 2022) in chat conversations, handling single-turn, in-context (Sun et al. 2024), and interleaved text and image inputs (Laurençon et al. 2024). MLLMs have benefited from vision-and-language tuning, enabling them to process diverse modalities (Wu et al. 2024). However, the general pretraining is insufficient for expert navigation tasks. Our work addresses this gap by adapting an MLLM from general scenarios to navigation through a specialized tuning approach.

Data Augmentation in Vision-and-Language Navigation

To overcome data scarcity in navigation, various data augmentation techniques (Zhao et al. 2021; Huang et al. 2019) have been proposed. These include utilizing a speaker module (Fried et al. 2018; Dou and Peng 2022) to generate synthetic instructions, leveraging multilingual data (Li, Tan, and Bansal 2022a), incorporating counterfactual information (Parvaneh et al. 2020; Fu et al. 2020), and altering environments (Li, Tan, and Bansal 2022b; Liu et al. 2021). For outdoor VLN, previous studies have explored training agents on instructions with different styles (Zhu et al. 2021), pretraining on auxiliary tasks (Armitage, Impett, and Senrich 2023) and using driving videos (Li et al. 2024). However, the use of auxiliary training data to tailor MLLMs for outdoor navigation remains largely unexplored. Our work addresses this gap by using synthetic data to adapt MLLMs for urban VLN tasks.

3 Method

We present FLAME (Flamingo-Architected Embodied Agent), an agent for urban Vision-and-Language Navigation (VLN). Our method comprises three key components: 1) An architecture for urban navigation. 2) A three-phase tuning technique. 3) A synthetic data creation process.

Task Formulation

We formalize urban VLN as follows: Given a navigation instruction $I = \{w_1, w_2, \dots, w_n\}$, an agent starts in an ini-

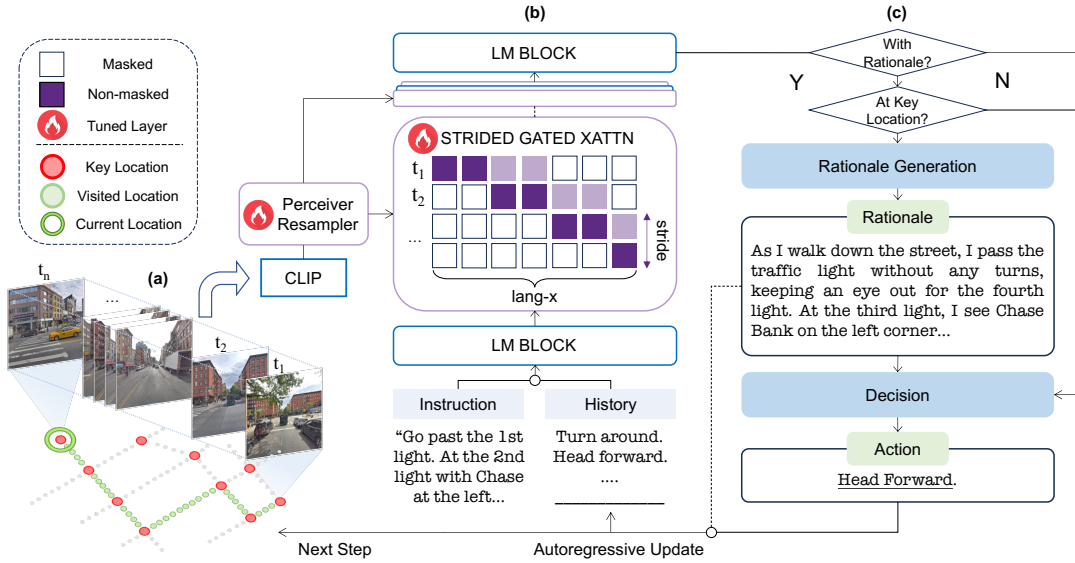


Figure 2: Overview of FLAME’s navigation process at time step t_n . The architecture, based on Flamingo, integrates vision modules for observation processing and decoder blocks for instruction and history handling. The finetuned STRIDED GATED XATTN layers prioritize recent observations in cross-attention computation. At key locations (intersections), FLAME can engage in reasoning before decision-making or proceed directly to action selection. The navigation process is autoregressive.

tial state S_0 . At each timestep t , the agent selects an action $a_t \in \mathcal{A}$ based on the current observation O_t and instruction I , where $\mathcal{A} = \{\text{Forward, Left, Right, Stop, Turn Around}\}$. The environment’s state transition function $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ updates the agent’s state to S_{t+1} . This process continues until the agent selects the Stop action. Navigation is successful if the agent stops at the target node or within one step of it. Figure 2(a) provides an illustration of agent’s observations and trajectory.

FLAME Architecture

FLAME builds upon the Flamingo architecture (Alayrac et al. 2022), leveraging cross-attention to process visual and textual inputs without extending context length, while ensuring single-pass fine-tuning per trajectory. We introduce two key adaptations to Flamingo for urban VLN:

Strided Cross-Attention To handle the large number of observations in urban VLN, we implement strided cross-attention (Child et al. 2019) in the cross attention layer, as depicted in Figure 2(b). The Perceiver Resampler module transforms CLIP (Radford et al. 2021) features into a compact set of visual tokens, with N_r tokens per observation. Let $X_t^v \in \mathbb{R}^{N \times d}$ represent flattened visual tokens up to timestep t , where $N = N_r \cdot t$. The LM block outputs a word embedding matrix X , with X_t as the the segment of context corresponds to t , and l as the stride length. Pattern $S = \{S_1, \dots, S_t\}$ defines visual token indices for attention at t , with $S_t = \{k, k+1, \dots, t \cdot N_r\}$ and $k = \max(0, (t-l) \cdot N_r)$. The strided cross-attention score is given by:

$$A(X_t, S_t) = \text{CMA}(X_t W_Q, (x_j^v W_K)_{j \in S_t}, (x_j^v W_V)_{j \in S_t}), \quad (1)$$

where CMA denotes cross-modal attention, W_Q, W_K^v, W_V^v are learnable parameters, and x_j^v refers to the j -th visual token. This approach is aimed at prioritizing recent observations, thereby augmenting the system’s proficiency in identifying significant features in the dynamic environment.

Action Prediction As Figure 2(c) shows, FLAME predicts actions based on the instruction I , current observation O_t , history of observations $O_{\leq t-1}$ and previous actions. By default, the next action a_t is obtained by:

$$a_t = \text{MLLM}(I, O_1, a_1, \dots, O_{t-1}, a_{t-1}, O_t). \quad (2)$$

Optionally, to make agent’s thought process transparent and understandable, FLAME generate rationales at key locations (i.e., intersections) before action prediction, similar to ReAct (Yao et al. 2023). Let R_t denote the generated rationale, the next action can be obtained by:

$$a_t = \text{MLLM}(I, O_1, a_1, \dots, O_{t-1}, a_{t-1}, O_t, R_t). \quad (3)$$

Three-Phase Tuning for Navigation

To adapt Flamingo for urban VLN tasks, we propose a three-phase tuning paradigm, as depicted in Figure 3(a).

Single Perception Tuning In the first phase, we train the model on a street view captioning task to strengthen its feature recognition abilities. Given a dataset $\mathcal{D}_{p1} = \{\tau^{(i)}\}_{i=1}^N$, where $\tau^{(i)} = \{(P^{(i)}, O^{(i)}, c^{(i)})\}$ consists of a caption prompt $P^{(i)}$, an observation $O^{(i)}$, and a ground truth caption $c^{(i)}$, the training objective is:

$$\mathcal{L}_{p1} = - \sum_{i=1}^N \log p(c^{(i)} | P^{(i)}, O^{(i)}; \theta), \quad (4)$$

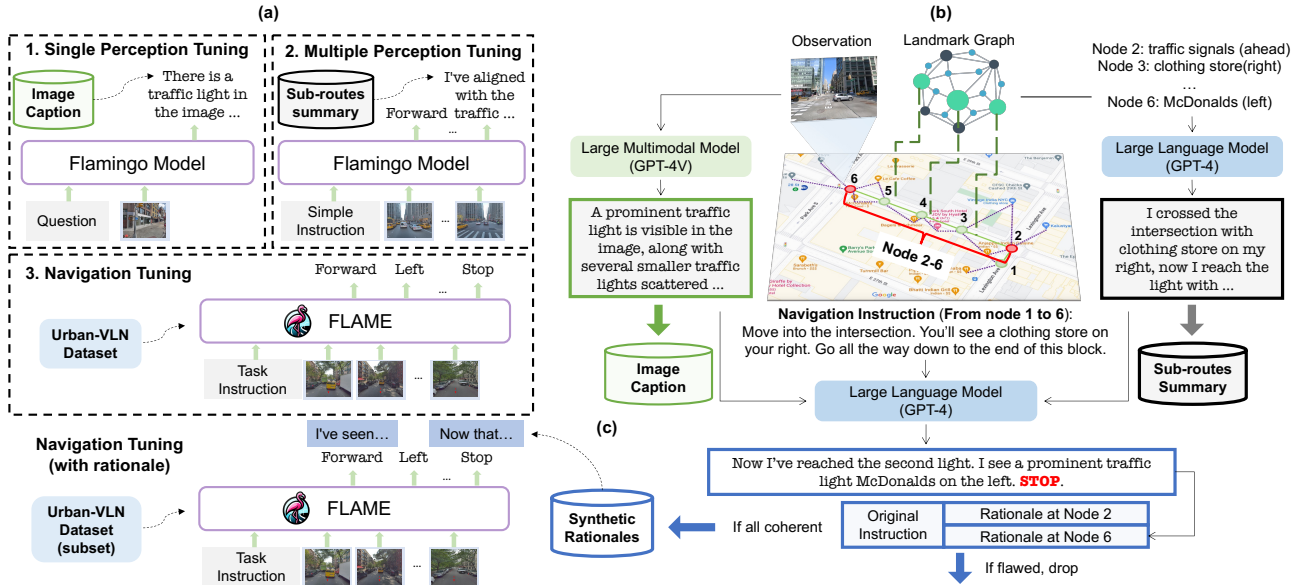


Figure 3: Illustration of the three-phase tuning for navigation and synthetic data generation process. (a) The first phase trains the model on single-perception tasks. The second phase escalates to handling multi-perceptual input. Finally, the model undergoes an end-to-end finetuning. (b) We utilize LLMs to generate street view captions, route summaries and simple instructions to aid the training of the first two phase. (c) We further synthesize rationales to validate the reasoning capability of FLAME.

where θ represents learnable parameters of the model.

Multiple Perception Tuning The second phase focuses on synthesizing sequential observations and performing actions. Using a dataset $\mathcal{D}_{p2} = \{\tau^{(i)}\}_{i=1}^N$, where each instance $\tau^{(i)} = \{(I^{(i)}, O_1^{(i)}, a_1^{(i)}, \dots, O_{T^{(i)}}^{(i)}, a_{T^{(i)}}^{(i)}, s^{(i)})\}$ consists of a simple instruction $I^{(i)}$, observations $O_{\leq T^{(i)}}$, actions $a_{\leq T^{(i)}}$ and a ground truth route summary $s^{(i)}$. We supervise two objectives:

$$\mathcal{L}_{sum}^{(i)} = \log p(s^{(i)} | I^{(i)}, O_{\leq T^{(i)}}^{(i)}; \theta), \quad (5)$$

$$\mathcal{L}_{act}^{(i)} = \sum_{t=1}^{T^{(i)}} \log p(a_t^{(i)} | I^{(i)}, O_{\leq t}^{(i)}, a_{\leq t-1}^{(i)}; \theta), \quad (6)$$

where $\mathcal{L}_{sum}^{(i)}$ is the route summarization loss and $\mathcal{L}_{act}^{(i)}$ is the imitation loss. The total loss for the second phase is:

$$\mathcal{L}_{p2} = - \sum_{i=1}^N (\mathcal{L}_{act}^{(i)} + \mathcal{L}_{sum}^{(i)}). \quad (7)$$

End-to-End Navigation Tuning Finally, FLAME is finetuned on a VLN dataset $\mathcal{D}_{nav} = \{\tau^{(i)}\}_{i=1}^N$, where $\tau^{(i)} = \{(I^{(i)}, O_1^{(i)}, a_1^{(i)}, \dots, O_{T^{(i)}}^{(i)}, a_{T^{(i)}}^{(i)})\}$. We minimize:

$$\mathcal{L}_{nav} = - \sum_{i=1}^N \sum_{t=1}^{T^{(i)}} \log p(a_t^{(i)} | I^{(i)}, O_{\leq t}^{(i)}, a_{\leq t-1}^{(i)}; \theta). \quad (8)$$

This multi-phased approach progressively builds the model’s capabilities, from environment understanding to complex decision-making.

Synthetic Data Generation

To support the finetuning of our agent, we leverage LLMs to automatically synthesize street view captions, route summaries (Figure 3(b)), and navigation rationales (Figure 3(c)).

Street View Caption Generation We focus on generating captions for street views at key locations, which are crucial for navigation. These images are processed by GPT-4V (Achiam et al. 2023) to produce descriptions. To ensure diversity, we randomly select meta-prompts from a hand-crafted prompt list, guiding the caption generation process. The prompt list contains specialized prompts refined manually to elicit diverse street view descriptions from GPT-4V.

Route Summary Generation We construct a comprehensive knowledge graph of landmarks by combining places of interests from OpenStreetMap with additional map information, following the methodology of (Schumann and Riezler 2021). Using this graph as input, we employ GPT-4 to generate detailed route summaries and simple navigation instructions between key locations.

Rationale Generation for VLN Datasets To validate reasoning capabilities of our agent, we generate synthetic rationales for the VLN datasets. The process involves segmenting each VLN trajectory into sub-routes, then retrieving the corresponding image caption and summary for each route. We then utilize GPT-4 to generate a synthetic rationale at each key location, based on the retrieved information. To maintain data quality, we discard instances with flawed rationales. This results in an altered subset of the VLN dataset augmented with synthetic rationales, which facilitates end-to-end training with reasoning capabilities.

Model	Touchdown						Map2seq					
	Dev Set			Test Set			Dev Set			Test Set		
	TC↑	SPD↓	nDTW↑	TC↑	SPD↓	nDTW↑	TC↑	SPD↓	nDTW↑	TC↑	SPD↓	nDTW↑
RCONCAT (2019)	10.60	20.4	22.50	11.80	20.40	22.90	17.10	-	30.70	14.70	-	27.70
GA (2019)	12.00	18.70	25.20	11.9	19.00	24.90	18.20	-	33.00	17.00	-	30.10
VLN-Trans (2021)	15.00	20.30	27.00	16.20	20.80	27.80	18.60	-	31.10	17.00	-	29.50
ARC+L2S (2020)	19.48	17.05	-	16.68	18.84	-	-	-	-	-	-	-
ORAR (2022)	30.05	11.12	45.50	29.60	11.79	45.30	49.88	5.87	62.70	47.75	6.53	62.10
VELMA (2024)	29.83	14.67	43.44	27.38	15.03	41.93	52.75	6.78	66.45	48.70	6.80	62.37
VLN-Video (2024)	34.50	9.60	-	31.70	11.2	-	-	-	-	-	-	-
Loc4Plan (2024)	34.50	10.50	-	32.90	11.50	-	48.00	7.00	-	45.30	7.20	-
RCONCAT* (2019)	9.27	19.14	16.40	8.34	21.02	15.70	8.25	38.71	12.51	7.63	41.22	12.03
GA* (2019)	7.12	20.91	13.27	6.63	21.76	12.44	6.50	38.74	11.67	6.07	41.25	10.88
ORAR* (2022)	20.99	10.97	30.78	21.34	11.02	32.44	49.88	6.74	62.96	48.53	7.22	61.91
FLAME	41.28	9.14	55.96	40.20	9.53	54.56	56.95	5.95	71.36	52.44	5.91	67.72

Table 1: Comparison with state-of-the-art models on Touchdown and Map2seq datasets. Models denoted by (*) utilize visual context consistent with FLAME’s implementation. **Bold** values indicate best performance.

4 Experiments

Experiment Setup

Datasets We evaluate our approach on two urban Vision-and-Language Navigation (VLN) datasets: Touchdown (Chen et al. 2019) and Map2seq (Schumann and Riezler 2021), both set in the StreetLearn environment (Mirowski et al. 2018). Touchdown contains 9,326 instruction-trajectory pairs, while Map2seq comprises 7,672 pairs. The augmented dataset for the first two training phases contains 2,354 and 4,674 instances, respectively. Our agent is benchmarked against others on the original datasets.

We collect 6,518 (out of 9,326) and 6,291 (out of 7,672) pairs grounded with rationales at key locations using GPT-4 for Touchdown and Map2seq, respectively. We formulate synthetic pairs to create a dataset for evaluating reasoning capabilities. The reasoning performance is evaluated exclusively on the subset containing synthetic rationales to ensure a fair comparison.

Metrics For the original VLN datasets, we employ three metrics for performance evaluation: Task Completion (TC), Shortest-Path Distance (SPD) and Normalized Dynamic Time Warping (nDTW). Specifically, TC represents the percentage of success. SPD calculates the minimum distance from the stop location to the goal. nDTW assesses the overlap between the agent’s and the ground truth trajectories.

To further evaluate the agent’s reasoning capabilities with synthetic rationales, we introduce two new metrics:

- Rationale Coherence (RC):

$$RC = \frac{\sum_{i=1}^N \sum_{j=1}^{M_i} \text{CFR}_{rc}(I_i, \gamma(R_i^j), R_i^j)}{\sum_{i=1}^N M_i} \quad (9)$$

- Rationale-Action Alignment (RA):

$$RA = \frac{\sum_{i=1}^N \sum_{j=1}^{K_i} \text{CFR}_{ra}(R_i^j, A_i^j)}{\sum_{i=1}^N K_i} \quad (10)$$

Here, I_i is the i -th instruction, R_i^j is the agent’s rationale, A_i^j is the action, and $\gamma(R_i^j)$ is the ground truth rationale at the j -th key location. M_i and K_i are the number of key location overlaps and visited key locations, respectively. N is the total number of instances. Functions $\text{CFR}_{rc}(\cdot)$ and $\text{CFR}_{ra}(\cdot)$ use GPT-4 (Achiam et al. 2023) to evaluate rationale consistency and action alignment (true or false), respectively.

Implementation Details Our agent is built upon Otter and OpenFlamingo (Li et al. 2023a; Awadalla et al. 2023), integrating CLIP (Radford et al. 2021) and LLaMA (Touvron et al. 2023). To accommodate CLIP’s input size, we crop and resize panoramas, which differs from the broader visual context used in existing models but aligns better with MLLM’s nature. In the Touchdown task, we randomize the agent’s heading at the start of each trajectory by selecting one of the possible directions based on its neighboring nodes. Training for the first two phases takes 1 hour each, while the navigation fine-tuning requires 12 hours on a single A100 GPU.

Comparison with SOTAs

In the section, we compare FLAME with previous state-of-the-art (SOTA) approaches. As shown in Table 1, FLAME establishes new SOTA performance on both Touchdown and Map2seq datasets. On Touchdown test split, FLAME surpasses the previous SOTA, Loc4Plan (Tian et al. 2024), by 7.3% in TC and 1.97% in SPD. This demonstrates the superiority of our MLLM-based approach in comprehending navigational instructions and environmental cues, resulting in higher success rates and better path adherence. For Map2seq, FLAME outperforms VELMA (Schumann et al. 2024), an LLM-based method, with a 3.74% increase in TC and a 5.35% improvement in nDTW. This highlights the advantage of Multimodal LLMs in capturing comprehensive information compared to text-only LLM.

Furthermore, we evaluated open-sourced methods in our visual setting, which features a restricted field of view similar to human vision. The results, marked with an asterisk (*)

Param	Touchdown (subset)						Map2seq (subset)					
	Dev Set			Test Set			Dev Set			Test Set		
	TC↑	RC↑	RA↑	TC↑	RC↑	RA↑	TC↑	RC↑	RA↑	TC↑	RC↑	RA↑
$T = 0.0, P = 1$	37.45	84.73	97.41	37.73	85.96	99.44	49.42	82.54	97.28	45.60	84.38	96.61
$T = 0.7, P = 4$	38.27	82.07	96.62	36.77	82.37	97.40	50.19	83.97	97.07	48.92	85.86	96.21
$T = 1.0, P = 4$	37.27	80.90	97.47	36.96	84.12	97.77	46.33	81.34	96.63	46.58	86.97	97.74
$T = 0.7, P = 8$	37.18	84.83	98.60	37.54	83.71	97.71	52.32	82.38	96.92	45.60	85.32	97.97
$T = 1.0, P = 8$	39.17	86.48	97.48	37.82	86.80	96.91	48.84	86.61	96.98	49.51	88.00	99.25

Table 2: Performance comparison of the agent’s reasoning capability on urban VLN dataset subsets under varying decoding temperatures and paths. The temperature T regulates the randomness in sampling, and the number of decoding paths P refer to the number of times the agent samples rationale-to-action pairs, with the final action determined through a voting process.

Method	Touchdown		Map2seq	
	RC	RA	RC	RA
Human	82.88	100.00	85.29	97.64
Vanilla	92.79	99.10	88.82	95.88
Calibrated	81.08	100.00	84.12	97.06

Table 3: Comparison of metric calculation methods.

in Table 1, show a significant performance drop for baseline models, especially on the Touchdown dataset. This underscores the dependence of baseline methods on panoramic environmental perception. In contrast, FLAME excels at navigating with non-panoramic visual input.

Reasoning Performance

We evaluated FLAME’s reasoning capabilities during navigation using the self-consistency approach (Wang et al. 2023), exploring various decoding paths and temperatures. The results are presented in Table 2. Rationale Coherence (RC) and Rationale-Action Alignment (RA) consistently remained above 80% and 95% respectively, indicating robust rationale generation and strong consistency. Higher temperatures led to performance fluctuations, especially with fewer decoding paths. However, when we increase the number of decoding paths to 8, we see more pronounced improvements in both TC and RC. On Touchdown, FLAME achieved optimal TC performance with a temperature of 1.0 and 8 decoding paths, outperforming greedy decoding ($T=0.0, P=1$) by 1.72% in TC and 1.75% in RC. For Map2seq, with 8 decoding paths, FLAME surpassed greedy decoding by 3.91% in TC and 3.62% in RC on the test set. These results demonstrate that increased sampling diversity and a larger decoding budget enable FLAME to generate diverse rationales and effectively ensemble reasoning results, leading to improved decision-making. This provides strong evidence for the reasoning capabilities developed through synthetic rationale tuning and the advanced decision-making capacity of FLAME’s architecture.

Metric Calculation To validate the reliability of our automatic metric calculations (Eqs. 9 and 10), we conducted human evaluations on 50 instances each from the Touchdown

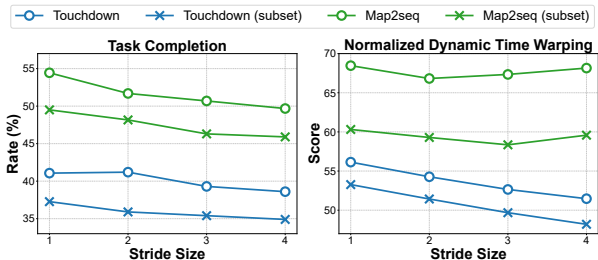


Figure 4: The effect of varying strides on TC and nDTW.

and Map2seq datasets. The results, presented in Table 3, revealed discrepancies between the vanilla method and human assessments. However, our calibrated approach significantly reduced these disparities, demonstrating the human-comparable reliability of the automatic evaluation method. Consequently, we employ this calibrated evaluation technique as default. Specifically, We calibrate the outputs of CFR_{rc} and CFR_{ra} as follows:

- When action at key location is incorrect: If $CFR_{ra} = 1$ but $CFR_{rc} = 1$, we force CFR_{rc} to 0.
- When action at key location is correct: If $CFR_{rc} = 1$ but $CFR_{ra} = 0$, we force CFR_{ra} to 1.

Analyses

This section presents a comprehensive analysis of our approach, evaluating the strided cross-attention module, the three-phase tuning technique, and providing qualitative insights of navigation details.

Effect of Strided Cross Attention To investigate the effect of the strided cross attention module, we conduct experiments on different stride sizes on both original datasets and subsets in Figure 4. We observe that increasing stride size generally correlates with a decrease in task completion rates, highlighting the importance of prioritizing current observations over longer history in decision-making processes. Paying full attention to numerous observations appears to overwhelm the agent, negatively impacting performance. Based on these findings, we set the default stride size to 1 in our implementation. However, nDTW scores for Map2seq fluctuate

Acknowledgments

This work was supported in part by the Natural Science Foundation of China under Grant 62303307, in part by National Key R&D Program of China under Grant No.2023YFB4705700, in part by Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102 and in part by the Fundamental Research Funds for the Central Universities.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Alteschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. GPT-4 Technical Report. arXiv:2303.08774.
- Alayrac, J. B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, 23716–23736.
- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and Van Den Hengel, A. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3674–3683.
- Armitage, J.; Impett, L.; and Sennrich, R. 2023. A priority map for vision-and-language navigation with trajectory plans and feature-location cues. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1094–1103.
- Awadalla, A.; Gao, I.; Gardner, J.; Hessel, J.; Hanafy, Y.; Zhu, W.; Marathe, K.; Bitton, Y.; Gadre, S.; Sagawa, S.; et al. 2023. OpenFlamingo: An Open-Source Framework for Training Large Autoregressive Vision-Language Models. arXiv:2308.01390.
- Chen, H.; Suhr, A.; Misra, D.; Snavely, N.; and Artzi, Y. 2019. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12538–12547.
- Chen, J.; Lin, B.; Xu, R.; Chai, Z.; Liang, X.; and Wong, K. Y. K. 2024. 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*, 9796–9810.
- Child, R.; Gray, S.; Radford, A.; and Sutskever, I. 2019. Generating Long Sequences with Sparse Transformers. arXiv:1904.10509.
- Dai, W.; Li, J.; Li, D.; Tiong, A.; Zhao, J.; Wang, W.; Li, B.; Fung, P.; and Hoi, S. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *Advances in Neural Information Processing Systems*, 49250–49267.
- Dou, Z. Y.; and Peng, N. 2022. Foam: A follower-aware speaker model For vision-and-language navigation. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4332–4340.
- Fried, D.; Hu, R.; Cirik, V.; Rohrbach, A.; Andreas, J.; Morency, L. P.; Berg Kirkpatrick, T.; Saenko, K.; Klein, D.; and Darrell, T. 2018. Speaker-follower models for vision-and-language navigation. In *Advances in Neural Information Processing Systems*.
- Fu, T. J.; Wang, X. E.; Peterson, M. F.; Grafton, S. T.; Eckstein, M. P.; and Wang, W. Y. 2020. Counterfactual vision-and-language navigation via adversarial path sampler. In *European Conference on Computer Vision*, 71–86.
- Huang, H.; Jain, V.; Mehta, H.; Ku, A.; Magalhaes, G.; Baldrige, J.; and Ie, E. 2019. Transferable representation learning in vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7404–7413.
- Ku, A.; Anderson, P.; Patel, R.; Ie, E.; and Baldrige, J. 2020. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 4392–4412.
- Laurençon, H.; Saulnier, L.; Tronchon, L.; Bekman, S.; Singh, A.; Lozhkov, A.; Wang, T.; Karamcheti, S.; Rush, A.; Kiela, D.; et al. 2024. Obelics: An open web-scale filtered dataset of interleaved image-text documents. In *Advances in Neural Information Processing Systems*, 71683–71702.
- Li, B.; Zhang, Y.; Chen, L.; Wang, J.; Yang, J.; and Liu, Z. 2023a. Otter: A Multi-Modal Model with In-Context Instruction Tuning. arXiv:2305.03726.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023b. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*, 19730–19742.
- Li, J.; Padmakumar, A.; Sukhatme, G.; and Bansal, M. 2024. Vln-video: Utilizing driving videos for outdoor vision-and-language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 18517–18526.
- Li, J.; Tan, H.; and Bansal, M. 2022a. Clear: Improving vision-language navigation with cross-lingual, environment-agnostic representations. In *Findings of the Association for Computational Linguistics: NAACL*, 633–649.
- Li, J.; Tan, H.; and Bansal, M. 2022b. Envedit: Environment editing for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15407–15417.
- Lin, B.; Nie, Y.; Wei, Z.; Zhu, Y.; Xu, H.; Ma, S.; Liu, J.; and Liang, X. 2024. Correctable landmark discovery via large models for vision-language navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 8534–8548.
- Liu, C.; Zhu, F.; Chang, X.; Liang, X.; Ge, Z.; and Shen, Y. D. 2021. Vision-language navigation with random environmental mixup. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1644–1654.

- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, 34892–34916.
- Long, Y.; Li, X.; Cai, W.; and Dong, H. 2024. Discuss before moving: Visual language navigation via multi-expert discussions. In *IEEE International Conference on Robotics and Automation*, 17380–17387.
- Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K. W.; Zhu, S. C.; Tafjord, O.; Clark, P.; and Kalyan, A. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *Advances in Neural Information Processing Systems*, 2507–2521.
- Mirowski, P.; Grimes, M.; Malinowski, M.; Hermann, K. M.; Anderson, K.; Teplyashin, D.; Simonyan, K.; Zisserman, A.; Hadsell, R.; et al. 2018. Learning to navigate in cities without a map. In *Advances in Neural Information Processing Systems*.
- Parvaneh, A.; Abbasnejad, E.; Teney, D.; Shi, J. Q.; and Van den Hengel, A. 2020. Counterfactual vision-and-language navigation: Unravelling the unseen. In *Advances in Neural Information Processing Systems*, 5296–5307.
- Qi, Y.; Wu, Q.; Anderson, P.; Wang, X.; Wang, W. Y.; Shen, C.; and Hengel, A. v. d. 2020. Reverie: Remote embodied visual referring expression in real indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9982–9991.
- Qiao, Y.; Qi, Y.; Yu, Z.; Liu, J.; and Wu, Q. 2023. March in chat: Interactive prompting for remote embodied referring expression. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15758–15767.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, 8748–8763.
- Schumann, R.; and Riezler, S. 2021. Generating landmark navigation instructions from maps as a graph-to-text problem. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 489–502.
- Schumann, R.; and Riezler, S. 2022. Analyzing generalization of vision and language navigation to unseen outdoor areas. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 7519–7532.
- Schumann, R.; Zhu, W.; Feng, W.; Fu, T.-J.; Riezler, S.; and Wang, W. Y. 2024. Velma: Verbalization embodiment of llm agents for vision and language navigation in street view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 18924–18933.
- Sun, Q.; Cui, Y.; Zhang, X.; Zhang, F.; Yu, Q.; Wang, Y.; Rao, Y.; Liu, J.; Huang, T.; and Wang, X. 2024. Generative multimodal models are in-context learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14398–14409.
- Tian, H.; Meng, J.; Zheng, W. S.; Li, Y. M.; Yan, J.; and Zhang, Y. 2024. Loc4plan: Locating before planning for outdoor vision and language navigation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 4073–4081.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M. A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. LLaMA: Open and Efficient Foundation Language Models. arXiv:2302.13971.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q. V.; Chi, E. H.; Narang, S.; Chowdhery, A.; and Zhou, D. 2023. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, 24824–24837.
- Wu, S.; Fei, H.; Qu, L.; Ji, W.; and Chua, T.-S. 2024. NExT-GPT: Any-to-Any Multimodal LLM. In *Proceedings of the 41st International Conference on Machine Learning*, 53366–53397.
- Xiang, J.; Wang, X.; and Wang, W. Y. 2020. Learning to stop: A simple yet effective approach to urban vision-language navigation. In *Findings of the Association for Computational Linguistics: EMNLP*, 699–707.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K. R.; and Cao, Y. 2023. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.
- Zhang, J.; Wang, K.; Xu, R.; Zhou, G.; Hong, Y.; Fang, X.; Wu, Q.; Zhang, Z.; and He, W. 2024. Navid: Video-based vlm plans the next step for vision-and-language navigation. In *Proceedings of Robotics: Science and Systems*.
- Zhao, M.; Anderson, P.; Jain, V.; Wang, S.; Ku, A.; Baldrige, J.; and Ie, E. 2021. On the evaluation of vision-and-language navigation instructions. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 1302–1316.
- Zheng, D.; Huang, S.; Zhao, L.; Zhong, Y.; and Wang, L. 2024. Towards learning a generalist model for embodied navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13624–13634.
- Zhou, G.; Hong, Y.; Wang, Z.; Wang, X. E.; and Wu, Q. 2024. Navgpt-2: Unleashing navigational reasoning capability for large vision-language models. In *European Conference on Computer Vision*, 260–278.
- Zhou, G.; Hong, Y.; and Wu, Q. 2024. Navgpt: Explicit reasoning in vision-and-language navigation with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7641–7649.
- Zhu, W.; Wang, X.; Fu, T. J.; Yan, A.; Narayana, P.; Sone, K.; Basu, S.; and Wang, W. Y. 2021. Multimodal text style transfer for outdoor vision-and-language navigation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 1207–1221.