

FloVerse: Floor Plan-Guided Multi-Modal Navigation

WeiQi Huang¹, Shuangyi Dong¹, Jiaxin Li¹, Yifei Guo¹, Zan Wang¹, Wei Liang¹✉

¹ School of Computer Science & Technology, Beijing Institute of Technology

<https://wikiahuang.github.io/floverse/>

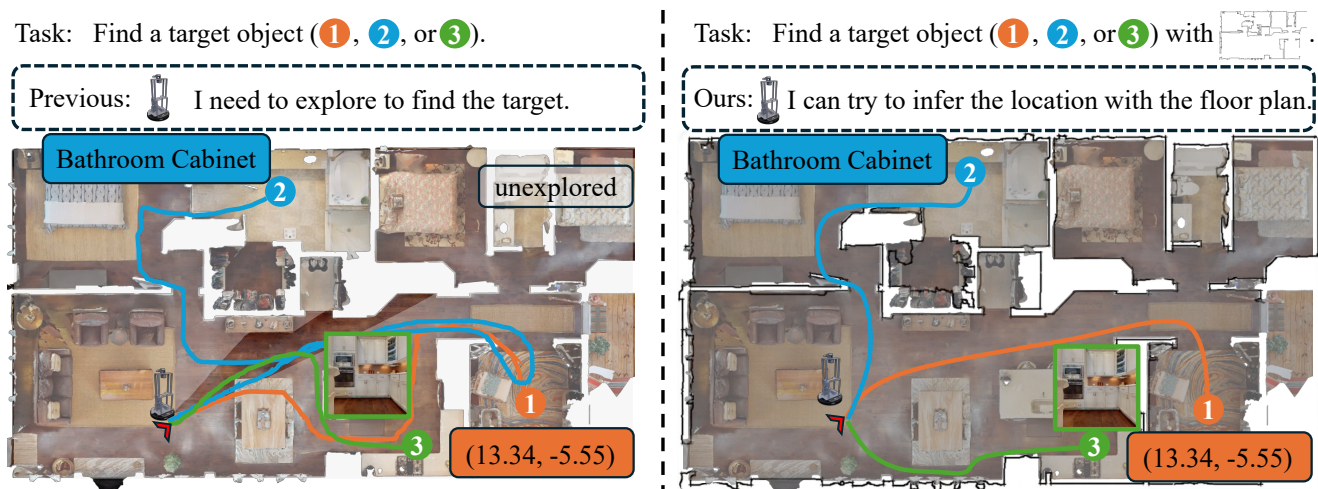


Figure 1. **Overview of FloVerse.** The agent is given a goal of any modality (PointNav, ObjectNav, or ImageNav). Prior methods rely only on local observations and must explore unseen areas. In contrast, **ThreeDiff** leverages floor-plan spatial priors for more efficient navigation.

Abstract

Floor plans encapsulate compact spatial priors, enabling agents to navigate unseen scenes more efficiently. While prior work has explored floor plan-guided navigation, it has focused mainly on PointNav and a limited set of environments. To bridge this gap, we introduce **FloVerse**, a new task for floor plan-guided embodied navigation that unifies PointNav, ObjectNav, and ImageNav. To support this **FloVerse**, we assemble **FloVerse-1.6K**, a large-scale dataset of 1.6K scenes from HM3D and Gibson 4+, paired with corresponding floor plans, comprising 240K expert trajectories and 12M RGBD frames. We further propose **ThreeDiff**, a two-stage imitation learning policy comprising a **planner**, a diffusion-based multimodal goal-reasoning module trained via masked-modality modeling, and a **refiner**, a depth-based trajectory-refinement module for safe execution. Extensive experiments demonstrate that (1) floor-plan priors improve navigation performance across all goal modalities, and (2) **ThreeDiff** implicitly captures spatial information from floor

plans. These results underscore the effectiveness of spatial priors and validate our proposed unified approach for floor plan-guided embodied navigation.

1. Introduction

Embodied navigation is advancing at an unprecedented pace, achieving impressive performance across challenging tasks such as point-goal (PointNav) [1, 25, 32, 41], object-goal (ObjectNav) [5, 14, 19, 36, 38], and image-goal (ImageNav) [16, 28, 29] navigation. Despite these achievements, most methods still depend on active mapping or extensive exploration, which often leads to short-sighted behaviors that reduce overall efficiency (Fig. 1).

Recent studies [9, 18] explore using floor plans as spatial priors to guide navigation [9, 18]. Floor plans are easily accessible and encode informative geometric structure, making them inherently well-suited for navigation in unseen environments. However, existing studies focus mainly on PointNav and limited scenes, leaving the broader potential of floor plan-guided navigation largely unexplored.

In fact, floor plans contain not only geometric structures

✉ Corresponding author: Wei Liang (liangwei@bit.edu.cn).

but also implicit semantic information about the environment layout, such as room functionalities and typical object distributions. These geometric and semantic regularities may provide rich spatial priors that enable agents to infer goal locations and plan more effectively, potentially enhancing embodied navigation across diverse goal modalities.

To investigate how floor plan priors can facilitate embodied navigation, we introduce **FloVerse**, a new task in which an agent must navigate to a goal—specified as a point, object, or image—using its egocentric RGBD observations and a floor plan in unseen environments (Fig. 1). To support this task, we develop a fully automated pipeline that derives floor plans from 3D scene meshes, producing a diverse collection of 1,627 floor plans across HM3D [22] and Gibson 4+ [26, 33] datasets. Note that each floor plan corresponds to a single level of a scene. We then leverage SpatialLM [20] to enrich object annotations in HM3D scenes, enhancing both scene coverage and category diversity. After manual verification, the resulting dataset contains 325 object categories across 299 scenes. Finally, we collect expert navigation trajectories across all three modalities, yielding 240K waypoint sequences and over 12M RGBD–pose pairs. Detailed statistics and the collection pipeline are provided in Sec. 3.

Building on this task, we propose **ThreeDiff**, a unified two-stage end-to-end policy that exploits both global floor plan spatial priors and local geometric cues to efficiently navigate across diverse goal modalities. In the first stage, a diffusion-based **planner** predicts a goal-conditioned coarse trajectory, capturing high-level spatial intent and long-horizon dependencies. In the second stage, the **refiner**, a local refinement module, utilizes depth-derived occupancy cues to adapt the coarse trajectory to the agent’s immediate surroundings, generating collision-free motions.

We extensively evaluate **ThreeDiff** on **FloVerse** to assess the impact of incorporating floor plan priors. Experimental results show that floor plan information consistently improves navigation efficiency and success rates across all goal modalities. Beyond quantitative gains, **ThreeDiff** exhibits an emergent ability to infer goals’ location, even without explicit supervision. Compared to specialized baselines trained for individual goal modalities, **ThreeDiff** delivers comparable or superior performance, highlighting the effectiveness and generality of the proposed unified framework.

Our main contributions are:

- **A multi-modal floor plan–guided navigation task.** **FloVerse** introduces navigation to point, object, or image goals using floor plans and egocentric observations.
- **A supporting dataset, FloVerse-1.6K.** It offers a rich set of reconstructed floor plans and expert navigation trajectories spanning a diverse range of scenes.
- **A two-stage navigation model.** **ThreeDiff** integrates global planning with local obstacle awareness, achieving strong performance.

- **Empirical validation of floor plan priors.** Quantitative and qualitative analysis show that incorporating floor plans consistently improves navigation, highlighting their value as structured spatial priors.

2. Related Work

2.1. Visual Navigation

Current visual navigation methods generally fall into two categories: mapping-based and mapless approaches. The former typically integrates semantic information into various map representations. Normally, researchers extract semantic categories [6, 11, 23, 40], pixel-level features [13, 38], and implicit image features [8, 21, 31] to construct comprehensive maps or topological graphs that support subsequent planning. In contrast, mapless visual navigation methods focus on directly translating observations into actions. Previous work has demonstrated notable success in PointNav through large-scale trial-and-error reinforcement learning [32]. Nomad [28] utilizes a diffusion policy to generate smooth trajectories for exploration and image-goal navigation. More recently, an increasing number of studies have leveraged the reasoning capabilities of large language models (LLMs) to guide agents in navigation [4, 7, 34, 42, 43]. However, mapping-based methods incur significant overhead from explicit map construction, whereas mapless methods lack global guidance when the goal is not observable. Both limitations hinder navigation efficiency.

2.2. Floor Plan Guided Navigation

Given the inherent limitations of the aforementioned methods, recent research has begun to explore floor plan-guided visual navigation. In indoor environments, floor plans provide readily available spatial priors that facilitate efficient visual navigation. However, the absence of detailed information about movable objects, such as furniture, makes floor plans temporally stable but also presents challenges for directly planning collision-free trajectories. To address this limitation, prior studies have incorporated Voronoi diagrams [27], pre-collected visual features [17], traversability cues from robot observations [9], and 2D point cloud data [10] to enhance floor plan representations and improve localization and navigation. FloNa [18] makes the first attempt to navigate in unseen environments using only a monocular camera and a floor plan within an end-to-end framework.

Although floor plans have been shown to improve navigation efficiency, prior evaluations have been limited to small-scale environments and the PointNav task. Their impact on other goal modalities, such as ImageNav and ObjectNav, is still largely unexplored. In this work, we scale up the evaluation and systematically examine whether floor plans can enhance navigation performance across PointNav, ObjectNav, and ImageNav.




Figure 2. **Examples of constructed floor plans.** Top row: generated floor plans. Bottom row: corresponding 3D scenes from the HM3D and Gibson 4+ datasets

2.3. Multi-Modal Goal-Oriented Navigation

Early navigation models typically address a single goal modality. More recently, research has shifted toward developing unified frameworks capable of handling multiple goal types within a single model. For instance, VIENNA [30] unifies four navigation tasks through an attention-based parse-and-query framework, demonstrating that a multi-task agent can match or surpass task-specific models. Goat-Bench [15] introduces a multi-modal lifelong navigation setting, where an agent processes goals expressed as images, object categories, and textual descriptions using a modular skill-chain architecture. Uni-Goal [37] represents all goal modalities as a unified graph structure and leverages an online scene graph to infer goal locations dynamically. NavDP [3] adopts a diffusion-based formulation, encoding goal types as modality tokens fused with RGBD inputs for joint policy learning. OmniNav [35] features a fast-slow collaborative architecture and large-scale multi-task training to achieve strong generalization across diverse navigation tasks. Building on these advances, we introduce the first model that unifies three goal modalities—PointNav, ObjectNav, and ImageNav—within a floor plan-guided navigation framework.

3. Dataset

3.1. Floor Plan Construction

To support **FloVerse**, we extract stable structural elements, such as walls, from object-containing meshes, as these elements correspond to features represented in floor plans. Vertices are first filtered based on their height and normal orientation, with a threshold of 1.25m, effectively removing most objects. To further mitigate reconstruction artifacts and spurious noise, we apply morphological operations, including dilation and erosion, for denoising. Using this pipeline, we construct a total of 1,627 floor plans, comprising 1,488 from the HM3D scenes and 139 from Gibson 4+ scenes. Examples of the resulting floor plans are shown in Fig. 2.

3.2. Object Goal Annotation

Most object-goal navigation methods involve a limited range of object categories, typically 6 [6] or 21 [2]. HM3D-OVON

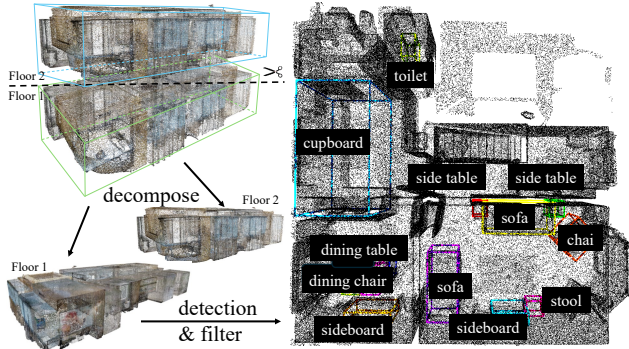


Figure 3. **Our object annotation pipeline.** Each scene point cloud is first decomposed into separate layers and then processed by SpatialLM for object detection. The detected objects are manually filtered to remove incorrect annotations, yielding the final results.

[39] substantially expands the coverage to 379 categories but includes only 181 scenes. We further clean this dataset by removing objects whose associated viewpoints fall outside navigable regions, resulting in 304 object categories across 169 scenes. To enrich category and scene diversity, we employ SpatialLM [20] for additional object recognition, and manually filter annotations to ensure quality. As illustrated in Fig. 3, SpatialLM is applied to each scene point cloud to detect objects and extract their bounding boxes and poses. For multi-layer scenes, we perform layer-wise decomposition to maintain accuracy. In total, **FloVerse-1.6K** contains 325 object categories spanning 299 scenes.

Table 1. **Data splits and the number of navigation episodes.** IO episodes refer to the episodes supporting ImageNav ObjectNav.

Dataset	train		eval	
	HM3D	Gibson 4+	HM3D	Gibson 4+
scenes	1321	121	167	18
total episodes	198,150	18,150	25,050	2,700
IO episodes	65,700	-	8,550	-

3.3. Trajectory Collection

We construct a large-scale set of expert trajectories covering all goal modalities for both imitation learning and evaluation. For PointNav, goals are randomly sampled positions on the navigable map (details of navigability are provided in the *supplementary material*), whereas for ObjectNav and ImageNav, goals correspond to annotated object locations. The start position is randomly chosen from navigable areas at least 5m away from the goal. The shortest path between each start-goal pair is computed using the A* algorithm and discretized into waypoints at 10cm intervals. At each waypoint, the agent’s pose and RGBD observations are recorded. Each expert trajectory includes RGBD-pose pairs, a floor plan, and a goal. The goal is defined as a point in PointNav, an image in ImageNav, and an object category in ObjectNav. As summarized in Tab. 1, the dataset contains about 240K

trajectories with 12M RGBD–pose pairs, including roughly 74K ImageNav/ObjectNav episodes (IO episodes).

4. Problem Definition

FloVerse aims to enable floor plan–guided embodied navigation in unseen indoor environments. Formally, given an environment \mathcal{E} that provides both perceptual observations and a global 2D floor plan F , the agent interacts with \mathcal{E} through a sequence of observations, actions, and state transitions. Specifically, at each time step t , the agent receives a state $\mathbf{s}_t = (I_t, D_t, \mathbf{p}_t)$, where $I_t \in \mathbb{R}^{H \times W \times 3}$ and $D_t \in \mathbb{R}^{H \times W}$ denote the egocentric RGB and depth images, and $\mathbf{p}_t = (x_t, y_t, \theta_t)$ represents the agent’s 2D pose in the environment. The floor plan $F \in \mathbb{R}^{H_f \times W_f}$ serves as a spatial prior describing the structural layout of \mathcal{E} . The goal g depends on the navigation modality:

$$g = \begin{cases} g_{\text{point}} \in \mathbb{R}^2, & \text{PointNav} \\ g_{\text{object}} \in \mathcal{L}, & \text{ObjectNav} \\ g_{\text{image}} \in \mathbb{R}^{H \times W \times 3}, & \text{ImageNav} \end{cases} \quad (1)$$

where \mathcal{L} is the set of textual object labels.

At each step, the agent executes an action $a_t \in \mathcal{A}$ and transitions according to

$$p_{t+1} = Tr(\mathcal{E}, p_t, a_t), \quad (2)$$

which updates its state and generates a new observation \mathbf{s}_{t+1} . The action space \mathcal{A} is either a discrete set of primitives $\mathcal{A}_{\text{discrete}} = \{\text{move forward}, \text{turn left}, \text{turn right}, \text{stop}\}$, or a continuous set of 2D waypoints $\mathcal{A}_{\text{continuous}} \subset \mathbb{R}^2$. An episode is represented as $\tau = (\mathbf{s}_t, a_t)_{t=1}^T$, where T is the episode length. An episode is considered successful if the agent reaches the goal g within a distance threshold d before exceeding the maximum number of steps T_{max} or collisions δ_{max} . The objective is to learn a policy $\pi(a_t | \mathbf{s}_t, F, g)$ that efficiently drives the agent to the goal while minimizing trajectory length and collisions.

5. Method

We introduce a two-stage navigation model that generates obstacle-aware trajectories for all three goal modalities using the agent’s egocentric observation and the floor plan. As shown in Fig. 4, the **planner** first predicts a coarse trajectory using the floor plan and goal information, where the goal modality is randomly masked during training to encourage modality-agnostic reasoning. The **refiner** then refines this trajectory using local geometric cues from the depth observation, producing a final path that avoids nearby obstacles.

5.1. Multi-Modality Goal Conditioned Planner

To enable a single model to support diverse goal modalities, we employ a random modality masking strategy during training. At each iteration, the model receives only one of the

three goal modalities, while the other two are masked. The resulting masked goal condition is defined as:

$$C_g = \mathbf{m} \circ (C_{g_{\text{point}}}, C_{g_{\text{object}}}, C_{g_{\text{image}}}) \quad (3)$$

where $\mathbf{m} = [m_{\text{point}}, m_{\text{object}}, m_{\text{image}}]$ is a binary vector with exactly one entry set to 1.

The point goal condition is encoded using an MLP:

$$C_{g_{\text{point}}} = \text{MLP}(g_{\text{point}}), \quad (4)$$

where $\text{MLP}(\cdot)$ denotes a multilayer perceptron. For ImageNav, we obtain the image-goal condition by encoding both the current RGB observation I_t and the goal image g_{image} with an image encoder $E(\cdot)$, instantiated as EfficientNet, followed by feature concatenation:

$$C_{g_{\text{image}}} = \text{MLP}(E(I_t) \oplus E(g_{\text{image}})). \quad (5)$$

Similarly, the object-goal condition is obtained using a CLIP encoder $\text{CLIP}(\cdot)$ applied to the current image and the target object label:

$$C_{g_{\text{object}}} = \text{MLP}(\text{CLIP}(I_t) \oplus \text{CLIP}(g_{\text{object}})). \quad (6)$$

This strategy provides two advantages. First, by exposing the model to all three modalities under the same input setting (s_t, F, \bar{a}_t) across training iterations, where \bar{a}_t represents the ground truth action, the modalities effectively inform and reinforce one another, leading to more generalizable trajectory generation. Second, by always combining the goal with features from the current observation, the model learns a consistent notion of current–goal correspondence across modalities, thereby improving stability and sample efficiency during diffusion training.

Then, we introduce a diffusion policy to learn a goal-oriented floor plan-guided navigation policy that takes C_g , $C_{\mathcal{F}}$, and $C_{\mathcal{O}}$ as inputs to generate the current action \hat{a}_t . Here, $C_{\mathcal{F}}$ denotes floor plan features extracted using EfficientNet, and $C_{\mathcal{O}}$ represents the current observation features, which is computed via a multi-head self-attention module (SA): $C_{\mathcal{O}} = SA(f_I^{t-l:t} \oplus f_D^{t-l:t} \oplus f_p^{t-l:t})$, where $f_I^{t-l:t}$, $f_D^{t-l:t}$, and $f_p^{t-l:t}$ represent the RGB, depth, and pose features from the past l timesteps, respectively.

To model $P_{\theta}(\bar{a}_t | C_g, C_{\mathcal{F}}, C_{\mathcal{O}})$, we deploy a conditional U-Net [24] and DDPM [12] scheduler to perform the forward and backward process. The training objective is:

$$\mathcal{L}_1 = \text{MSE}(\epsilon_k, \epsilon_{\theta}(\bar{a}_t + \epsilon_k, k)), \quad (7)$$

where MSE represents the mean squared error, ϵ denotes the noise label, and k represents the denoising steps.

5.2. Depth-Based Trajectory Refiner

The training trajectories are optimal, collision-free paths, meaning the model does not receive collision feedback,

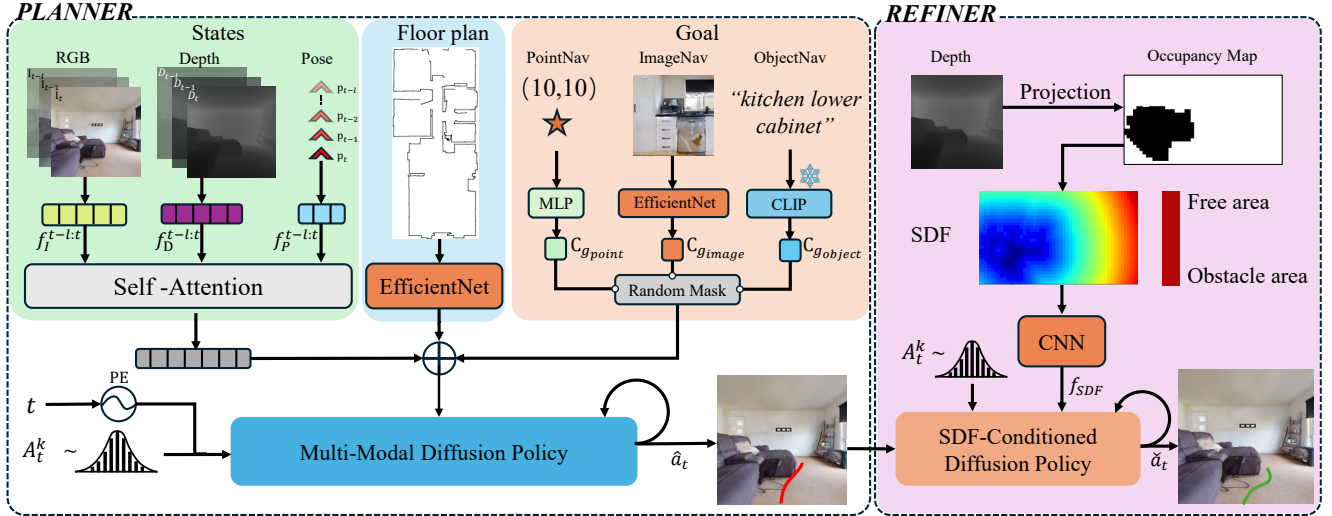


Figure 4. **Overview of ThreeDiff.** ThreeDiff is a two-stage trajectory generation framework for floor plan-guided navigation. In the first stage, goal-related features, floor plan representations, and observation embeddings are concatenated and used as the conditioning input to a diffusion model, which generates a coarse initial trajectory. In the second stage, depth-derived obstacle-awareness features are integrated with the initial trajectory and fed into a second diffusion model, which refines the trajectory to yield a final, obstacle-aware navigation path.

which limits its ability to actively avoid obstacles. Since obstacle avoidance is fundamentally tied to understanding spatial geometry, we address this limitation by training a second diffusion model for safer trajectory prediction using the geometric information from depth. As shown in Fig. 4, we first project the depth into a local binary occupancy map from a Bird’s Eye View (BEV) perspective, with a grid size of 20×50 and a resolution of 0.1m, where 0 represents obstacles and 1 represents free space. Next, we compute a local Signed Distance Field (SDF), which is then encoded using a CNN. The resulting embedding f_{SDF} is combined with the predicted trajectory \hat{a}_t and fed into the second diffusion model, which outputs the refined trajectory \tilde{a}_t . The diffusion process aims to fit the following distribution: $P_\phi(\tilde{a}_t|\hat{a}_t, f_{SDF})$.

In addition to minimizing the discrepancy between the predicted and true trajectories, \tilde{a}_t should also be as distant as possible from obstacles. Thus, the training objective is defined as:

$$\mathcal{L}_2 = \text{MSE}(\epsilon_k, \epsilon_\phi(\tilde{a}_t + \epsilon_k, k)) + \alpha \mathcal{L}_{\text{collision}}, \quad (8)$$

and $\mathcal{L}_{\text{collision}} = \frac{1}{N} \sum_{i=1}^N \exp(-\lambda \cdot \rho_i)$, where N is the number of waypoints, ρ_i is the l_2 distance from the i -th point to the nearest obstacle, where α and λ are hyperparameters that controls the weight of the collision loss. Note that, to ensure differentiability, ρ_i is computed from the SDF after applying linear interpolation.

5.3. Implement Details

In our implementation, we first train the first diffusion model until convergence, and then integrate the second diffusion model, training both jointly until the system reaches con-

vergence. For the object goal condition, we utilize a frozen CLIP-ViT-B/32 model to encode the images and text labels. All other image encoders in ThreeDiff are EfficientNet-B0 with non-shared weights, trained from scratch. The multi-head attention mechanism consists of four heads, each with four layers. We optimize the models using the AdamW optimizer, with the learning rate governed by PyTorch’s CosineAnnealingLR scheduler. The maximum learning rate is set to 0.0001, and training is conducted for a maximum of 20 epochs. The parameter α is set to 0.1, and λ is set to 1. The training of ThreeDiff is performed on 4 NVIDIA 4090 GPUs, with a batch size of 32 per GPU.

6. Experiments

In this section, we examine four key aspects: (1) the impact of floor plans on performance in PointNav, ImageNav, and ObjectNav (Sec. 6.2); (2) the cross-modality complementarity effect under random modality masking (Sec. 6.3); (3) the performance of existing navigation models on FloVerse (Sec. 6.4); and (4) the contribution of the depth-based refiner to the overall performance of ThreeDiff (Sec. 6.5).

6.1. Setup and Metrics

Setup Following [18], we conduct our experiments in the Gibson simulator [33]. To ensure consistent evaluation across different methods and mitigate potential failures caused by environment mesh imperfections, both discrete and continuous action outputs are converted into global positions and orientations for execution. At inference, ThreeDiff predicts the next 16 waypoints, and the first 10 are selected as the next action sequence. Additionally, to reduce stochasticity, ThreeDiff and its variants generate 30 trajectories per

Table 2. Comparison between ThreeDiff and ThreeDiff w/o F.

Method	PointNav		ImageNav		ObjectNav	
	SR	SPL	SR	SPL	SR	SPL
ThreeDiff w/o F	25.8	25.6	22.6	18.4	20.9	16.5
ThreeDiff	42.0	36.6	28.9	22.4	28.6	22.3

inference step and average them to produce the next action sequence. We employ a safety mechanism to mitigate collisions. Specifically, if a collision occurs, the agent returns to its previous state, reorients itself toward the goal, and then infers the following action based on the updated state. We record the collision events throughout the episode. An episode is considered successful if the agent reaches within a distance d of the goal within at most T_{\max} steps and incurs fewer than δ_{\max} collisions. In our experiments, we set $d = 1$, $T_{\max} = 500$, and $\delta_{\max} = 15$.

Metrics We evaluate performance using Success Rate (SR) and Success Weighted by Path Length (SPL). In ObjectNav, multiple instances of the same category may exist within a scene. An episode is deemed successful once the agent reaches any instance of the target category, while SPL is computed relative to the shortest geodesic path from the starting position to the specific instance encountered.

6.2. Impact of Floor-Plan Priors

To assess the floor plan’s impact across PointNav, ImageNav, and ObjectNav, we train a variant of the policy, denoted **ThreeDiff w/o F**, from scratch, which removes the floor plan from the input while keeping all other components identical to **ThreeDiff**. As shown in Tab. 2, removing the floor plan consistently reduces performance: SR drops by 16.2%, 6.3%, and 7.7%, and SPL decreases by 11.0%, 4.0%, and 5.8% for PointNav, ImageNav, and ObjectNav, respectively. Note that because Gibson 4+ lacks image and object goals (Tab. 1), ImageNav and ObjectNav are evaluated only in HM3D, whereas PointNav is evaluated in both Gibson 4+ and HM3D.

The consistent reduction in both SR and SPL indicates that **ThreeDiff** effectively leverages floor-plan spatial priors to achieve more reliable and more efficient navigation. The most significant improvement appears in PointNav, where **ThreeDiff** outperforms **ThreeDiff w/o F** by 16.2% in SR and 11.0% in SPL. This substantial gain arises because PointNav benefits directly from the relatively explicit geometric priors provided by the floor plan, which guide the agent toward a precise goal location in previously unseen environments and significantly reduce navigation uncertainty. In contrast, ImageNav and ObjectNav show minor improvements. For these tasks, the floor plan offers only coarse semantic cues—such as room layout and spatial organization—rather than explicit object- or image-level information. While the policy can make limited use of these cues, their indirect nature limits their contribution to the required semantic rea-

Table 3. Comparison between ThreeDiff trained on individual modalities and jointly trained with random modality masking.

Method	PointNav		ImageNav		ObjectNav	
	SR	SPL	SR	SPL	SR	SPL
Point-only	42.1	38.9	-	-	-	-
Image-only	-	-	25.4	20.7	-	-
Object-only	-	-	-	-	24.3	19.9
ThreeDiff	42.0	36.6	28.9	22.4	28.6	22.3

soning, yielding more modest gains than in PointNav.

As depicted in Fig. 5, we conduct additional qualitative experiments to examine how floor-plan priors guide planning in ImageNav and ObjectNav. We visualize the mean initial predicted trajectories in four scenes—00123-C3ifY177Ldq, 00176-wcJYziD5pmF, 00744-1S7LAXRdDqK, and 00037-oKFJo8jzRW. In each scene, an object is selected as the target, its rendered view is used as the image goal, and multiple starting points are randomly sampled. Across scenes, **ThreeDiff** tends to produce trajectories that move more directly toward the goal, while those from **ThreeDiff w/o F** exhibit greater variance and more frequent exploratory movements. This contrast indicates that incorporating floor plan priors—such as coarse room layout cues—may help the agent better infer goal locations and improve navigation efficiency in ImageNav and ObjectNav.

6.3. Cross-Modality Complementarity

To investigate the complementarity among different goal modalities, we conduct a comparative ablation study in which three single-modality policies—Point-only, Image-only, and Object-only—are trained separately, while **ThreeDiff** is trained with all modalities using random masking. As shown in Tab. 3, **ThreeDiff** performs comparably with the Point-only model on PointNav but consistently outperforms all single-modality models on ImageNav and ObjectNav.

Interestingly, the Point-only policy achieves even marginally higher PointNav performance than **ThreeDiff**, and it substantially outperforms the Image-only and Object-only models by over 15% in SR and nearly 20% in SPL. This pattern suggests that explicit geometric cues in the floor plan are easier to learn.

The improved ImageNav and ObjectNav results of **ThreeDiff** further imply that the spatial knowledge learned from PointNav generalizes to semantic goal modalities, providing beneficial structural cues such as room connectivity and coarse spatial layout. In contrast, the slight drop in **ThreeDiff** on PointNav indicates that the additional training signals from ImageNav and ObjectNav introduce more implicit, layout-level semantic biases in the encoder, which slightly interfere with the learning of explicit geometric structures required by PointNav.

Finally, Fig. 6 presents the evolution of cosine similarity between predicted and ground-truth trajectories for different



Figure 5. **Qualitative results demonstrating the guiding effect of the floor plan.** Top row: Across four HM3D scenes, we define a goal location and multiple starting points. The initial trajectories of **ThreeDiff** (solid lines) and its variant without floor plan input (dashed lines) are shown under **image-goal** and **object-goal** conditions.




Figure 6. **Comparison of trajectory similarity changes during training.** Higher similarity indicates better model fitting ability.

training strategies. The curves show that **ThreeDiff** and the Point-only model maintain similarly high alignment, both clearly outperforming Image-only and Object-only. This result again reflects that PointNav supervision drives the encoder to capture explicit geometric priors, whereas ImageNav and ObjectNav supervision encourage more implicit semantic layout understanding. Given that all models share the same floor-plan encoder, these results collectively highlight how different modalities emphasize complementary aspects of spatial representation.

6.4. Comparison with Existing Models

To provide a comprehensive and fair evaluation of **ThreeDiff**, we compare it with several representative navigation methods. Specifically, we evaluate **DD-PPO** [32], an RL model pretrained on Gibson, using its depth variant with a ResNet50 + LSTM512 backbone, without any additional fine-tuning. We also assess **ZSON** [19], an ObjectNav model trained on HM3D, directly using its released weights. In addition, we benchmark two imitation learning approaches—**Nomad** [28] and **FloDiff** [18]. Since both are originally RGB-only, we augment them with the same depth encoder used in **ThreeD-**

Table 4. **Performance of baselines on the PointNav task in the Gibson 4+ and HM3D scenes.**

Method	Gibson 4+		HM3D	
	SR	SPL	SR	SPL
DD-PPO [32]	18.3	10.5	13.7	7.0
FloDiff (pretrain)[18]	40.0	28.8	23.7	17.8
FloDiff (finetune)[18]	38.8	29.6	27.7	21.1
ThreeDiff	54.4	50.0	38.1	32.5

iff and fine-tune their pre-trained checkpoints on our dataset to ensure a fair comparison. Furthermore, we include three multi-modal navigation models from [15], enabling a direct comparison with existing multi-modality baselines.

As shown in Tab. 4, DD-PPO achieves an 18.3% success rate in the Gibson 4+ scenes (compared to 95.6% SPL on the Habitat Challenge 2019 [26] validation set), despite being trained on all of these scenes. This drop is mainly due to task differences: unlike prior simpler settings that ignore collisions with scene objects, our task emphasizes obstacle avoidance and safer navigation—capabilities that DD-PPO does not possess. FloDiff (pre-trained) achieves success rates of 40% in Gibson 4+ and 23.7% in HM3D, owing to its integration of floor plan information, which enhances its planning capabilities. After fine-tuning, FloDiff shows a slight drop on Gibson 4+ but gains 4% on HM3D, indicating that incorporating depth information can further improve performance. **ThreeDiff** attains the best overall results, thanks to two key factors: (1) it has learned richer associations between multiple modalities and floor plans, leading to stronger spatial reasoning and navigation planning, and (2) it exhibits enhanced obstacle awareness, benefiting from the second-stage trajectory refinement (see Sec. 6.5).

As shown in Tab. 5, on ImageNav, **ThreeDiff** outperforms Nomad (both pretrained and fine-tuned) in success rate and SPL, achieving 6.1% higher SR and 5.4% higher SPL than

Table 5. Evaluation on the ImageNav and ObjectNav tasks.

Method	ImageNav		ObjectNav	
	SR	SPL	SR	SPL
ZSON [19]	-	-	7.2	1.1
NoMad (pretrain) [28]	15.4	10.9	-	-
NoMad (finetune) [28]	22.8	17.0	-	-
ThreeDiff	28.9	22.4	28.6	22.3

Table 6. Results of ThreeDiff on the GOAT-Bench [15]. We compare ThreeDiff with the three official baselines.

Method	ImageNav		ObjectNav	
	SR	SPL	SR	SPL
RL Monolithic [15]	11.7	7.3	25.7	13.7
RL Skill Chain[15]	42.2	18.0	25.8	17.0
Modular Goat[15]	27.9	19.5	29.4	17.0
ThreeDiff	31.0	25.6	34.5	30.8

the fine-tuned version, despite both being trained on the same expert trajectories. This indicates that incorporating floor-plan priors facilitates more reliable and efficient image navigation. On ObjectNav, ZSON performs poorly, mainly because its training episodes are short (under 5m), while most episodes in our test set exceed 7m, limiting its ability to handle long-range goal inference.

We further compare **ThreeDiff** with RL Monolithic, RL Skill Chain, and Modular GOAT [15]. To enable direct comparison, we evaluate **ThreeDiff** on the val-seen split of GOAT-Bench [15]. As shown in Tab. 6, on ImageNav, **ThreeDiff** attains a lower success rate than RL Skill Chain but achieves a 7.6% higher SPL, indicating more efficient path planning. On ObjectNav, **ThreeDiff** outperforms all three baselines in both success rate and SPL, demonstrating that leveraging floor-plan spatial information facilitates more reliable and efficient navigation.

It is worth noting that Modular GOAT, a map-based approach, achieves higher SPL than both RL Monolithic and RL Skill Chain, but still falls short of **ThreeDiff** by 6.1% on ImageNav and 13.8% on ObjectNav. While explicit mapping aids in preserving environmental memory and supports long-horizon planning, it requires additional movements during initial exploration, which reduces overall efficiency. In contrast, **ThreeDiff** exploits spatial priors in the floor plan, including both geometric structures and semantic regularities, providing guidance that reduces unnecessary exploration and enables more reliable and efficient navigation across tasks.

6.5. Ablation of Refiner

To assess the impact of the refiner, we ablate the second-stage module, producing the w/o refiner variant. As reported in Tab. 7, w/o refiner incurs performance degradation across all tasks, with the most significant decline observed on Point-

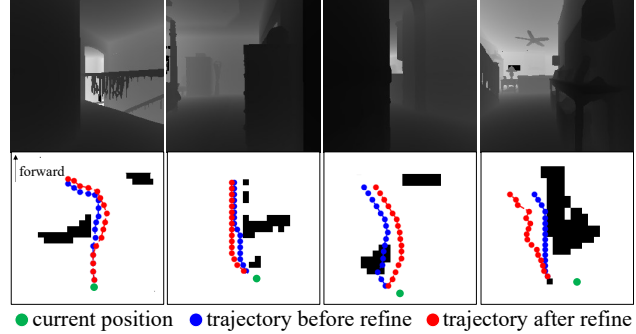


Figure 7. Impact of the second-stage refinement. Top row: depth observation at an inference step. Bottom row: corresponding local occupancy map with trajectories before and after refinement.

Table 7. Performance of w/o refiner and ThreeDiff.

Method	PointNav		ImageNav		ObjectNav	
	SR	SPL	SR	SPL	SR	SPL
w/o refiner	34.3	29.2	27.7	21.4	25.3	19.4
ThreeDiff	42.0	36.4	28.9	22.4	28.6	22.3

Nav. This underscores the critical role of the second-stage refinement in enhancing overall trajectory quality. Fig. 7 provides a visual comparison of trajectories before and after refinement. The refined trajectory clearly avoids obstacles, demonstrating that by converting depth information into local SDF, the model effectively captures navigability cues and generates safer waypoints.

7. Conclusion

We present **FloVerse**, a novel task that challenges agents to perform multimodal goal-directed navigation using floor plans and egocentric RGBD observations. To support research on this task, we collect **FloVerse-1.6K**, a large-scale dataset suitable for both training and evaluation. Furthermore, we introduce **ThreeDiff**, a two-stage end-to-end model capable of effectively handling multiple navigation modalities. Comprehensive quantitative and qualitative assessments demonstrate (1) floor plans significantly improve navigation efficiency, and (2) **ThreeDiff** achieves robust performance across all three navigation tasks. We believe our work will serve as a solid foundation for future advances in visual navigation.

Limitations **ThreeDiff** has two limitations. First, it has been primarily evaluated on floor plans from residential indoor environments, and its performance on other types of layouts remains to be explored. Second, it has not yet been tested in real-world scenarios, leaving its practical deployment to be investigated in future work.

Acknowledgments This project is supported by Beijing Natural Science Foundation (L242094) and the National Natural Science Foundation of China (NSFC, 62172043).

References

- [1] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018. **1**
- [2] Dhruv Batra, Aaron Gokaslan, Aniruddha Kembhavi, Oleksandr Maksymets, Roozbeh Mottaghi, Manolis Savva, Alexander Toshev, and Erik Wijmans. ObjectNav Revisited: On Evaluation of Embodied Agents Navigating to Objects. In *arXiv:2006.13171*, 2020. **3**
- [3] Wenzhe Cai, Jiaqi Peng, Yuqiang Yang, Yujian Zhang, Meng Wei, Hanqing Wang, Yilun Chen, Tai Wang, and Jiangmiao Pang. Navdp: Learning sim-to-real navigation diffusion policy with privileged information guidance. *arXiv preprint arXiv:2505.08712*, 2025. **3**
- [4] Yihan Cao, Jiazhao Zhang, Zhinan Yu, Shuzhen Liu, Zheng Qin, Qin Zou, Bo Du, and Kai Xu. Cognav: Cognitive process modeling for object goal navigation with llms. In *International Conference on Computer Vision (ICCV)*, 2025. **2**
- [5] Devendra Singh Chaplot, Dhiraj Gandhi, Abhinav Gupta, and Ruslan Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. **1**
- [6] Devendra Singh Chaplot, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. **2, 3**
- [7] 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 *AAAI Conference on Artificial Intelligence (AAAI)*, 2025. **2**
- [8] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Think global, act local: Dual-scale graph transformer for vision-and-language navigation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. **2**
- [9] Zu Lin Ewe, Fu-Hao Chang, Yi-Shiang Huang, and Li-Chen Fu. Spatial graph-based localization and navigation on scale-less floorplan. *IEEE Robotics and Automation Letters (RA-L)*, 9(4):3932–3939, 2024. **1, 2**
- [10] RG Goswami, H Sinha, PV Amith, J Hari, P Krishnamurthy, J Rizzo, and F Khorrami. Floor plan based active global localization and navigation aid for persons with blindness and low vision. *IEEE Robotics and Automation Letters (RA-L)*, 9(12):11058–11065, 2024. **2**
- [11] Yu Guo, Jinsheng Sun, Ruiheng Zhang, Zhiqi Jiang, Zhenqiang Mi, Chao Yao, Xiaojuan Ban, and Mohammad S Obaidat. An object-driven navigation strategy based on active perception and semantic association. *IEEE Robotics and Automation Letters (RA-L)*, 9(8):7110–7117, 2024. **2**
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. **4**
- [13] Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. In *International Conference on Robotics and Automation (ICRA)*, 2023. **2**
- [14] Hao Huang, Yu Hao, Congcong Wen, Anthony Tzes, Yi Fang, et al. Gamap: Zero-shot object goal navigation with multi-scale geometric-affordance guidance. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. **1**
- [15] Mukul Khanna, Ram Ramrakhya, Gunjan Chhablani, Sriram Yenamandra, Theophile Gervet, Matthew Chang, Zsolt Kira, Devendra Singh Chaplot, Dhruv Batra, and Roozbeh Mottaghi. Goat-bench: A benchmark for multi-modal lifelong navigation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. **3, 7, 8, A4**
- [16] Jacob Krantz, Theophile Gervet, Karmesh Yadav, Austin Wang, Chris Paxton, Roozbeh Mottaghi, Dhruv Batra, Jitendra Malik, Stefan Lee, and Devendra Singh Chaplot. Navigating to objects specified by images. In *International Conference on Computer Vision (ICCV)*, 2023. **1**
- [17] Jun Li, Chee Leong Chan, Jian Le Chan, Zhengguo Li, Kong Wah Wan, and Wei Yun Yau. Cognitive navigation for indoor environment using floorplan. In *International Conference on Intelligent Robots and Systems (IROS)*, 2021. **2**
- [18] Jiaxin Li, Weiqi Huang, Zan Wang, Wei Liang, Huijun Di, and Feng Liu. Flona: Floor plan guided embodied visual navigation. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2025. **1, 2, 5, 7**
- [19] Arjun Majumdar, Gunjan Aggarwal, Bhavika Devnani, Judy Hoffman, and Dhruv Batra. Zson: Zero-shot object-goal navigation using multimodal goal embeddings. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. **1, 7, 8, A4**
- [20] Yongsen Mao, Junhao Zhong, Chuan Fang, Jia Zheng, Rui Tang, Hao Zhu, Ping Tan, and Zihan Zhou. Spatialllm: Training large language models for structured indoor modeling. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. **2, 3, A1**
- [21] Yiyuan Pan, Yunzhe Xu, Zhe Liu, and Hesheng Wang. Planning from imagination: Episodic simulation and episodic memory for vision-and-language navigation. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2025. **2**
- [22] Santhosh K Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alex Clegg, John Turner, Eric Underlander, Wojciech Galuba, Andrew Westbury, Angel X Chang, et al. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In *Proceedings of the Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*, 2021. **2**
- [23] Santhosh Kumar Ramakrishnan, Devendra Singh Chaplot, Ziad Al-Halah, Jitendra Malik, and Kristen Grauman. Poni: Potential functions for objectgoal navigation with interaction-free learning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. **2**
- [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015. **4**

- [25] Pascal Roth, Julian Nubert, Fan Yang, Mayank Mittal, and Marco Hutter. Viplanner: Visual semantic imperative learning for local navigation. In *International Conference on Robotics and Automation (ICRA)*, 2024. 1
- [26] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 7
- [27] Vachirasuk Setalaphruk, Atsushi Ueno, Izuru Kume, Yasuyuki Kono, and Masatsugu Kidode. Robot navigation in corridor environments using a sketch floor map. In *International Symposium on Computational Intelligence in Robotics and Automation.*, 2003. 2
- [28] Ajay Sridhar, Dhruv Shah, Catherine Glossop, and Sergey Levine. Nomad: Goal masked diffusion policies for navigation and exploration. In *International Conference on Robotics and Automation (ICRA)*, 2023. 1, 2, 7, 8, A4
- [29] Xinyu Sun, Peihao Chen, Jugang Fan, Jian Chen, Thomas Li, and Mingkui Tan. Fgprompt: fine-grained goal prompting for image-goal navigation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 1
- [30] Hanqing Wang, Wei Liang, Luc Van Gool, and Wenguan Wang. Towards versatile embodied navigation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 3
- [31] Hanqing Wang, Wei Liang, Luc Van Gool, and Wenguan Wang. Dreamwalker: Mental planning for continuous vision-language navigation. In *International Conference on Computer Vision (ICCV)*, 2023. 2
- [32] Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *International Conference on Learning Representations (ICLR)*, 2020. 1, 2, 7
- [33] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 5
- [34] Tianyu Xu, Jiawei Chen, Jiazhao Zhang, Wenyao Zhang, Zekun Qi, Minghan Li, Zhizheng Zhang, and He Wang. Mmnav: Multi-view vla model for robust visual navigation via multi-expert learning. *arXiv preprint arXiv:2510.03142*, 2025. 2
- [35] Xinda Xue, Junjun Hu, Minghua Luo, Shichao Xie, Jintao Chen, Zixun Xie, Kuichen Quan, Guo Wei, Mu Xu, and Zedong Chu. Omninav: A unified framework for prospective exploration and visual-language navigation. In *International Conference on Learning Representations (ICLR)*, 2026. 3
- [36] Hang Yin, Xiuwei Xu, Zhenyu Wu, Jie Zhou, and Jiwen Lu. Sg-nav: Online 3d scene graph prompting for llm-based zero-shot object navigation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 1
- [37] Hang Yin, Xiuwei Xu, Linqing Zhao, Ziwei Wang, Jie Zhou, and Jiwen Lu. Unigoal: Towards universal zero-shot goal-oriented navigation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [38] Naoki Yokoyama, Sehoon Ha, Dhruv Batra, Jiuguang Wang, and Bernadette Bucher. Vlfm: Vision-language frontier maps for zero-shot semantic navigation. In *International Conference on Robotics and Automation (ICRA)*, 2024. 1, 2
- [39] Naoki Yokoyama, Ram Ramrakhya, Abhishek Das, Dhruv Batra, and Sehoon Ha. Hm3d-ovon: A dataset and benchmark for open-vocabulary object goal navigation. In *International Conference on Intelligent Robots and Systems (IROS)*, 2024. 3
- [40] Xinyao Yu, Sixian Zhang, Xinhang Song, Xiaorong Qin, and Shuqiang Jiang. Trajectory diffusion for objectgoal navigation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [41] Xiaoming Zhao, Harsh Agrawal, Dhruv Batra, and Alexander G Schwing. The surprising effectiveness of visual odometry techniques for embodied pointgoal navigation. In *International Conference on Computer Vision (ICCV)*, 2021. 1
- [42] Duo Zheng, Shijia Huang, Lin Zhao, Yiwu Zhong, and Liwei Wang. Towards learning a generalist model for embodied navigation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [43] Gengze Zhou, Yicong Hong, and Qi Wu. Navgpt: Explicit reasoning in vision-and-language navigation with large language models. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2024. 2