

LoGoPlanner: Localization Grounded Navigation Policy with Metric-aware Visual Geometry

Jiaqi Peng^{*,1,2}, Wenzhe Cai^{*,2}, Yuqiang Yang^{*,2}, Tai Wang^{2,†}, Yuan Shen^{1,2,†} and Jiangmiao Pang²

Abstract—Trajectory planning in unstructured environments is a fundamental and challenging capability for mobile robots. Traditional modular pipelines suffer from latency and cascading errors across perception, localization, mapping, and planning modules. Recent end-to-end learning methods map raw visual observations directly to control signals or trajectories, promising greater performance and efficiency in open-world settings. However, most prior end-to-end approaches still rely on separate localization modules that depend on accurate sensor extrinsic calibration for self-state estimation, thereby limiting generalization across embodiments and environments. We introduce LoGoPlanner, a localization-grounded, end-to-end navigation framework that addresses these limitations by: (1) finetuning a long-horizon visual-geometry backbone to ground predictions with absolute metric scale, thereby providing implicit state estimation for accurate localization; (2) reconstructing surrounding scene geometry from historical observations to supply dense, fine-grained environmental awareness for reliable obstacle avoidance; and (3) conditioning the policy on implicit geometry bootstrapped by the aforementioned auxiliary tasks, thereby reducing error propagation. We evaluate LoGoPlanner in both simulation and real-world settings, where its fully end-to-end design reduces cumulative error while metric-aware geometry memory enhances planning consistency and obstacle avoidance, leading to more than a 27.3% improvement over oracle-localization baselines and strong generalization across embodiments and environments. The code and models have been made publicly available on the project page.

I. INTRODUCTION

Autonomous navigation, requiring robots to reliably reach specified goals in unstructured environments, remains a fundamental challenge in robotics. Traditional navigation pipelines are typically modular, decomposing the task into perception, localization, mapping, and planning [1], [2], [3]. While this factorization improves interpretability and allows for component-level optimization, it often introduces compounding latency and suffers from cascading errors between modules [4]. These issues become particularly acute in real deployments, such as legged robots, where gait-induced vibrations in both cameras and IMUs reduce the accuracy of odometry and mapping—which in turn destabilizes downstream planning [3].

End-to-end learning-based methods [5], [6], [7], [8], [9] have recently emerged as a promising alternative, offering compact pipelines that map raw visual observations directly

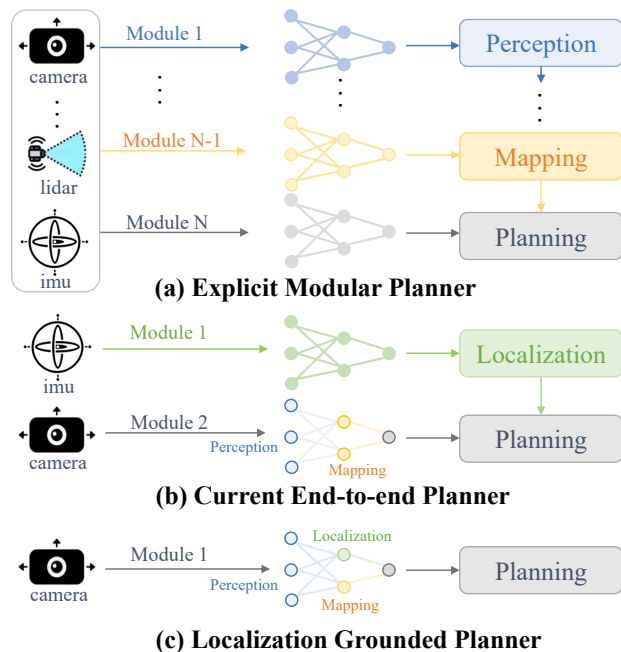


Fig. 1. (a) Traditional modular planners decompose tasks into modules, introducing cascading errors. (b) Existing end-to-end frameworks directly map observations to control signals but still rely on explicit localization modules. (c) LoGoPlanner integrates implicit state estimation and metric-aware geometry perception into policy for fully end-to-end planning.

to control signals or trajectories. Beyond mitigating cascading errors, such approaches also demonstrate high efficiency and ease of deployment in open-world scenarios. However, most approaches mainly replace perception, mapping, and planning modules but still rely on explicit localization modules such as SLAM or visual odometry [6], [9] for self-state updates, which requires precise extrinsic calibration between the camera and the robot chassis. This reliance arises because these planners typically process only single frame [5] or short clips [10], [9], lacking the ability to summarize long-term histories for consistent state updates. Without temporal grounding, short-term estimates accumulate errors over time, leaving trajectory planning vulnerable to drift and inconsistency. Also, single frame perception lacks geometric memory needed for robust metric reasoning [11]. Most methods reconstruct only partial or scale-ambiguous geometry, and fails to capture broader spatial context including occluded and rear-view regions, limiting the fidelity of spatial reasoning.

To address these challenges, we propose LoGoPlanner, a localization-grounded end-to-end navigation framework that integrates temporal visual geometry estimation with diffusion-based trajectory generation.

*Equal contribution.

†Corresponding author.

¹Jiaqi Peng, Yuan Shen are with the Department of Electronic Engineering, Tsinghua University.

²Jiaqi Peng, Wenzhe Cai, Yuqiang Yang, Tai Wang, Jiangmiao Pang and Yuan Shen are all with Shanghai AI laboratory.

This work is supported by Shanghai Artificial Intelligence Laboratory

Specifically, to overcome the reliance on explicit localization, we design an implicit state estimation module that operates on long visual sequences. Consecutive image histories are processed using pretrained visual-geometry backbones [12] for camera extrinsic estimation. We further decouple chassis and camera extrinsics: perception is tied to the camera viewpoint while control is executed at the chassis level. This separation prevents viewpoint-dependent perception errors from propagating into the control space, enabling stable action generation even under varying sensor placements. Training with camera data captured under varying heights and pitch angles enables cross-embodiment and cross-view ego-motion estimation. Additionally, implicit state features are used to project the goal back into the current frame, ensuring consistent goal alignment.

Furthermore, we aggregate multi-frame visual features and pose embeddings to generate an implicit reconstruction of the surrounding scene, thereby providing fine-grained geometric priors for planning. Visual geometry backbones such as VGGT [12] output relative-scale reconstructions, which cannot be directly aligned with planning trajectories. To resolve this, we introduce an efficient post-training procedure that incorporates depth information as a scene-scale prior, enabling the prediction of dense, metric-scale point clouds relative to the robot’s current position. Inspired by planning-oriented frameworks [13], we adopt a query-based policy architecture that fuses implicit state [14], geometry, and goal features into a unified planning context. Unlike approaches that pass explicit poses or point clouds downstream and thereby risk error accumulation, our query-driven design enables the diffusion head to operate directly on geometrically grounded representations. This allows the policy to iteratively refine trajectories toward feasible, collision-free solutions, with planning as the ultimate optimization target.

In simulation, LoGoPlanner achieves a 27.3% relative improvement over baselines that rely on oracle localization, demonstrating the effectiveness of implicit self-state estimation. Furthermore, in diverse real-world scenarios, our framework exhibits robust generalization across different embodiments and environments.

II. RELATED WORK

A. Learning-based planner

Recent end-to-end visual navigation frameworks aim to directly map visual inputs to control commands, eliminating the need for traditional modular pipelines. By leveraging semantic and geometric cues, these methods demonstrate the ability to plan across varied terrains. Some prior works adopt a supervised learning paradigm, where robots are trained to imitate expert trajectories or human demonstrations, reasoning about expert actions from vision-based observations [15], [16], [11], [17]. Beyond imitation learning, reinforcement learning has been applied to end-to-end navigation by optimizing policies through trial-and-error in simulation [8], [18], but its high sample complexity and sparse rewards make training costly. To overcome these challenges, Yang et al. [6] proposed iPlanner, which reformulates planning as an offline

bi-level optimization problem, improving efficiency but only relying on a single frame to capture the robot’s surrounding geometry. These methods are still trained in an open-loop manner, predicting the entire trajectory from start to goal. As a result, they lack explicit estimation of intermediate states and typically assume access to metric geometry, treating localization as an external input. Our work builds on these advances by further integrating state estimation and geometry understanding into navigation planning.

B. Video-geometric model

Recent progress in video-based geometry models [12], [19], [20] has significantly advanced multi-frame 3D scene understanding. For instance, Video Depth Anything [21] extends monocular depth estimation to long video sequences, exploiting temporal information to preserve geometric quality and consistency over time while retaining generalization. Beyond depth estimation, models such as VGGT [12] perform full video-based reconstruction. VGGT is a feed-forward neural network that jointly predicts dense 3D attributes, including depth maps, 3D point tracks, and camera extrinsics, from one or more views of a scene. By leveraging long temporal windows, these models provide geometrically consistent reconstructions and explicit camera pose estimates, offering fine-grained priors that are particularly valuable for downstream tasks such as environment perception and navigation [22].

C. Monocular Visual Odometry

Monocular visual odometry (VO) and SLAM methods inherently suffer from scale ambiguity. Classical geometric methods such as MonoSLAM [23] and PTAM [24] rely on handcrafted features and epipolar constraints but require strong motion or scene priors, and are prone to drift accumulation. Direct methods such as DSO [25] and LSD-SLAM [26] optimize photometric consistency, reducing reliance on features but still needing scale initialization and remaining sensitive to illumination changes. Learning-based methods introduce data-driven priors. Depth-assisted approaches such as MonoDepth [27] and ORB-SLAM2 [28] recover scale from monocularly estimated depth, but often generalize poorly across domains. End-to-end networks like PoseNet [29] regress pose and scale directly, while transformer-based models such as MonoRec [30] alleviate drift through long-range dependencies; however, both remain vulnerable to dynamic or fast-motion scenarios. Fusion-based approaches combine multiple cues: BEV-ODOM [31] exploits spatial regularities, CodedVO [32] improves scale estimation via coded apertures, and VINS-Mono [33] integrates IMU data to enhance robustness, albeit at the cost of additional sensors. Despite these advances, most monocular VO methods still depend on external priors—whether geometric assumptions, scale initialization, or auxiliary hardware—limiting their scalability and generalization.

III. PROBLEM FORMULATION

We study the problem of point-goal navigation using only RGB-D observations. An agent must navigate from

its start pose to a designated target point $g \in \mathbb{R}^3$ while avoiding collisions without relying on additional modules. Time proceeds in discrete steps $i = 1, \dots, N$. At each step, the agent receives an RGB-D observation $\mathbf{O}_i = (\mathbf{I}_i, \mathbf{D}_i)$, where $\mathbf{I}_i \in \mathbb{R}^{H \times W \times 3}$ is the RGB image and $\mathbf{D}_i \in \mathbb{R}^{H \times W}$ is the corresponding depth map. To successfully reach the goal, the agent must continuously estimate its own state from long-term visual history while simultaneously perceiving the surrounding environment to ensure safe navigation.

The predicted state of the agent up to time i is given by $\hat{s}_{1:i} = f(O_{1:i})$, which further enables the transformation of the global goal g into the current coordinate frame, yielding the relative goal $\hat{g}_i = f(\hat{s}_{1:i}, g)$. Based on this goal and the perceived environment, the agent plans an obstacle-avoiding trajectory $\tau_{1:T} = \{p_1, \dots, p_T\}$ toward the goal expressed in the chassis coordinate frame.

Unlike existing end-to-end approaches that depend on explicit localization modules, our framework performs implicit, closed-loop state estimation directly from visual sequences. At each step, the agent maintains (i) its estimated chassis pose \hat{p}_i in metric scale, (ii) the relative goal position \hat{g}_i in the current frame, and (iii) a dense local point cloud \hat{P}_i that captures its surrounding environment.

IV. METHODOLOGY

A. Overall Architecture

As illustrated in Fig. 2, LoGoPlanner is a unified end-to-end navigation framework that jointly learns metric-aware perception, implicit localization, and trajectory generation.

The system builds on a pretrained video geometry backbone augmented with depth-derived scale priors. Through auxiliary supervision on point prediction and pose estimation, the model encodes both fine-grained structures and long-term ego-motion into compact world point embeddings aligned with the planning coordinate system. A query-based design allows task-specific queries to extract state and geometry features through cross-attention and fuse them with goal embeddings into a unified planning context. A diffusion-based policy head then conditions on this context to iteratively refine noisy actions into collision-free trajectories.

B. Metric-aware Visual Geometry Learning

Recent video geometry models such as VGGT [12] can recover dense 3D scene geometry from image sequences, but their predictions are only defined up to an unknown scale, limiting their applicability to navigation. To address this limitation, we fine-tune the backbone by injecting scale priors from depth maps, enabling metric-scale scene reconstruction.

Concretely, given a causal sequence of N RGB images $\mathbf{I}_{1:N}$ from the same navigation scene, the video geometric model uses a vision transformer [34] to patchify each image \mathbf{I}_i into a set of K initial patch tokens $\mathbf{t}_i^I \in \mathbb{R}^{K \times C^I}$. The alternating-attention mechanism alternates between intra-frame and inter-frame attention, improving both local fidelity and long-horizon consistency. To inject scale priors into the semantic patches, we employ a lightweight variant of the same vision transformer [34] to encode input depth maps into

geometric tokens $\mathbf{t}_i^D \in \mathbb{R}^{K \times C^D}$, which are then fused at the patch level with initial semantic tokens \mathbf{t}_i^I . A transformer decoder module, integrating attention mechanisms with Rotary Position Embedding (RoPE) [35], is further used to integrate information within frames, producing per-frame features that are enriched with metric-scale information.

$$\mathbf{t}_i^{metric} = \text{Attention}(\text{RoPE}((\mathbf{t}_i^I, \mathbf{t}_i^D), \text{pos})), \quad (1)$$

where $\text{pos} \in \mathbb{R}^{K \times 2}$ denotes the spatial position coordinates of image patches, $\text{RoPE}(\cdot, \text{pos})$ applies position-dependent rotations to tokens using 2D coordinates in pos , thereby augmenting the intra-frame attention mechanism to better capture spatial positional relationships between patch tokens, and \mathbf{t}_i^{metric} representing the fused feature embedding at time i with metric-scale awareness.

To improve the accuracy of point cloud prediction, we introduce auxiliary tasks that provide additional supervision during training. Specifically, given the multi-frame feature representation $\mathcal{F} = \{\mathbf{t}_1^{metric}, \dots, \mathbf{t}_N^{metric}\}$ from the context RGB frames, we branch the features into two task-specific heads: a local point head and a camera pose head.

Local point prediction. For each frame i , the local point head ϕ_p maps metric-aware tokens \mathbf{t}_i^{metric} to a latent feature representation \mathbf{h}_i^p , which is further decoded to predict canonical local 3D points in the camera coordinate system:

$$\mathbf{h}_i^p = \phi_p(\mathbf{t}_i^{metric}), \quad \hat{P}_i^{\text{local}} = f_p(\mathbf{h}_i^p) \quad (2)$$

where $\hat{P}_i^{\text{local}} = \{\hat{\mathbf{p}}_{i,j}^{\text{local}}\}_{j=1}^M$ denotes a set of predicted local points j for frame i , which is supervised by local points in the camera coordinate system using the pinhole model for each pixel (u, v) in images:

$$\mathbf{p}_{\text{cam},i}(u, v) = D_i(u, v) K^{-1} [u \ v \ 1]^\top \quad (3)$$

Unlike directly using noisy depth maps for local point projection, our data-driven prediction improves reconstruction robustness and provides implicit features for world point prediction.

Camera pose prediction. In parallel, the camera pose head ϕ_c maps \mathbf{t}_i^{metric} to another task-specific feature \mathbf{h}_i^c , which is decoded into a predicted camera-to-world transformation:

$$\mathbf{h}_i^c = \phi_c(\mathbf{t}_i^{metric}), \quad \hat{T}_{c,i} = f_c(\mathbf{h}_i^c) \quad (4)$$

where the world coordinate system is defined with respect to the chassis frame of the last time step, ensuring consistency between the predicted camera trajectory and the planning coordinate system.

World point prediction. Rather than directly encoding the predicted local points or poses, we exploit their intermediate task-specific features to perform implicit transformations, thereby obtaining a compact representation of world points for downstream planning. For each frame i , we concatenate the local-point feature \mathbf{h}_i^p and the pose feature \mathbf{h}_i^c . These per-frame fused features are then aggregated across the patches by a context fusion module \mathcal{A} :

$$\mathbf{h}_i^w = \mathcal{A}([\mathbf{h}_i^p, \mathbf{h}_i^c]). \quad (5)$$

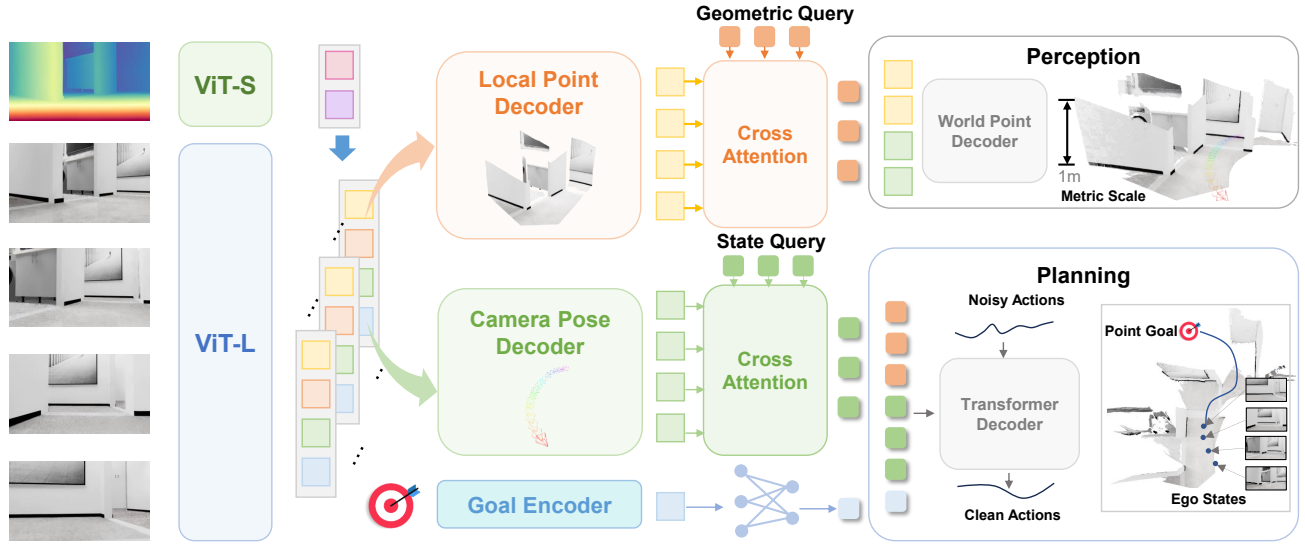


Fig. 2. **Architecture overview.** Our method injects scale priors into the image patches that are encoded by ViT [34], and finetunes the video geometry model to metric scale prediction. We adopt a query-based design in which ego state representation and environment geometry are implicitly aggregated through task-specific queries. A diffusion policy head is detached to generate feasible and collision-free trajectories.

Finally, the aggregated representation \mathbf{h}_i^w is passed through a point-cloud decoder ψ to upsample to the target resolution. The intermediate output \mathbf{z}_i is further processed to enforce a scale-invariant range and preserve sign information:

$$\mathbf{z}_i = \psi(\mathbf{h}_i^w), \quad \hat{P}_i^{\text{world}} = \text{sign}(\mathbf{z}_i) \cdot (\exp(|\mathbf{z}_i|) - 1) \quad (6)$$

where $\exp(|\cdot|) - 1$ is used to enhance expressiveness of large coordinate values while avoiding saturation, \hat{P}_i^{world} represents the reconstructed metric-scale scene point cloud expressed in the last frame to align the coordinate with planning trajectory.

C. Localization Grounded Navigation Policy

In navigation, perception is tied to the camera view-point, whereas control is executed at the chassis level. This mismatch leads to alignment errors, particularly in cross-embodiment settings where camera height and pitch may vary significantly. Traditional methods rely on explicit extrinsic calibration between the camera and chassis to align perception with planning, but such calibration is fragile and fails to generalize across different views.

In contrast, we decouple the estimation of camera and chassis poses into separate prediction tasks and leverage implicit feature interaction to bridge perception and planning without explicit calibration. Specifically, we assume the robot-mounted camera has no yaw rotation relative to the chassis. The model also predicts the chassis pose and latest goal position from pose estimation task-specific feature \mathbf{h}_i^c :

$$\hat{T}_{b,i} = f_b(\mathbf{h}_i^c) \quad \hat{g}_i = f_g(\mathbf{h}_i^c, g) \quad (7)$$

which is defined on the ground plane as (x_i, y_i, θ_i) , where (x_i, y_i) denotes the current position relative to the starting point, and θ_i is the yaw angle. Formally, the camera pose at time i is obtained from the chassis pose through the extrinsic transformation:

$$T_{b,i} = T_{c,i} \cdot T_{\text{ext}}, \quad (8)$$

where $T_{b,i}$ is the chassis transformation with respect to the start point, and T_{ext} denotes the fixed extrinsic transformation capturing the camera's height and pitch angle relative to the chassis. To endow the model with robustness across embodiments, we construct training data under arbitrary camera heights and varying pitch angles, thereby enabling the system to generalize across diverse camera configurations.

To achieve end-to-end planning, our approach does not explicitly feed the predicted extrinsics or point clouds into the network for goal transformation or trajectory optimization. Instead, inspired by UniAD [13], we adopt a **query-based design** in which different modules are aggregated through task-specific queries. Interactions across modules are realized via query cross attention. We set state queries Q_s to extract implicit state representation from pose specific tokens and geometric queries Q_g to extract implicit environment geometry from world point specific tokens, thus providing sufficient information for trajectory planning:

$$Q_S = \text{CrossAttn}(Q_s, \mathbf{h}^c) \quad (9)$$

$$Q_G = \text{CrossAttn}(Q_g, \mathbf{h}^p) \quad (10)$$

where the generated Q_S , Q_G and goal embedding are concatenated and passed to transformer decoder to produce the planning context query Q_P : These implicit features, which encode states and geometric properties, serve as conditioning signals for planning. This strategy avoids cascading errors that would otherwise arise from directly applying upstream predictions to downstream tasks, while ensuring that the final optimization target remains the trajectory planning error.

For navigation trajectory planning, we attach a diffusion policy head to generate action chunks $\{\mathbf{a}_t = (\Delta x_t, \Delta y_t, \Delta \theta_t)\}_{t=1}^T$. Starting from \mathbf{a}^K sampled from Gaussian noise, the model predicts noise from noisy action sequences, performs K iterative steps of denoising to produce a series of intermediate actions with decreasing levels of noise:

$$\mathbf{a}^{k-1} = \alpha(\mathbf{a}^k - \gamma \epsilon_\theta(Q_P, \mathbf{a}^k, k) + \mathcal{N}(0, \sigma^2 I)) \quad (11)$$

where \mathbf{a}^k denotes the noisy action at step k , ϵ_θ is the noise prediction network conditioned on planning context query Q_P , α and γ are the standard diffusion schedule parameters. This formulation enables the policy to iteratively refine actions toward feasible, collision-free trajectories.

V. EXPERIMENTS

A. Datasets and Implementation Details

We use a large-scale navigation dataset [36], generated with a simulation pipeline designed to efficiently generate diverse robot trajectories across a variety of 3D environments. The robot is modeled as a cylindrical rigid body with a differential-drive two-wheel configuration and is equipped with an RGB-D camera mounted on the top. To simulate embodiment variations across different robotic platforms, the robot’s height is randomized between 0.25 m and 1.25 m, while the camera’s pitch angle is randomized between 0° and 30° . Initial paths between randomly sampled start and goal points are generated using the A* algorithm. These paths are refined through greedy search and subsequently smoothed via cubic spline interpolation to ensure collision-free navigation. The dataset comprises over 200k trajectories and approximately 10 million rendered images.

We adopt a two-stage training paradigm. In the first stage, we fine-tune the decoder of the video geometry model and the task-specific head with a batch size of 12 for 24 hours. During this process, depth-based scale priors are injected, and supervision is provided by metric-scale scene point clouds and camera extrinsics. In the second stage, we jointly train the diffusion head together with the task-specific head while keeping the decoder of the backbone frozen. This stage uses a batch size of 32 and runs for three days, ensuring both robust state estimation and stable perception capabilities.

B. Main Results

We evaluate the performance of different learning-based planners in both simulation and real-world environments. The test environments are unseen during training, requiring the robot to continuously estimate its state, plan collision-free trajectories, and navigate toward the designated goal. Planning performance is measured using Success Rate (SR) and Success weighted by Path Length (SPL). The experimental results are summarized in Table I and Table II.

1) *Simulation Experiments*: To simulate realistic environments, we selected 40 scenes from the InternScenes dataset [37], including 20 home and 20 commercial scenes. Home scenes are characterized by narrow passages and cluttered semantic layouts, while commercial scenes cover representative categories such as hospitals, supermarkets, restaurants, schools, libraries, and offices. In each scene, 100 start–goal pairs are randomly sampled in unoccupied spaces with distances of 4–10 meters, and initial orientations are determined through path planning to avoid collisions.

Table I reports navigation performance for two scene categories. In the table, the “Localization” column indicates whether the planner has access to ground-truth localization



Fig. 3. Home scenes are characterized by narrow passages and cluttered semantic layouts, while commercial scenes cover representative categories such as hospitals, supermarkets, restaurants, schools, libraries, and offices.

TABLE I
SIMULATION RESULTS

Planner	Localization	Home		Commercial	
		SR \uparrow	SPL \uparrow	SR \uparrow	SPL \uparrow
DD-PPO [8]	×	–	–	–	–
	✓	0.4	0.4	5.3	5.2
iPlanner [6]	×	41.7	40.2	53.1	51.8
	✓	43.0	40.6	54.6	52.8
ViPlanner [5]	×	44.0	42.8	61.3	60.1
	✓	45.0	43.2	63.7	61.9
LoGoPlanner	×	57.3	52.4	67.1	63.9

× With explicit or implicit localization.

✓ With oracle localization from simulator.

from the simulator. A “✓” means the planner uses simulator-provided poses, while a “×” indicates no access to ground-truth localization: for iPlanner and ViPlanner, an external visual odometry module (ORB-SLAM3 [38] with RGB-D input) is used, whereas LoGoPlanner performs implicit state estimation without any external localization.

Reinforcement-learning-based planners like DD-PPO [8] typically require massive interaction data and careful reward shaping. They can underperform in unseen environments due to overfitting to training distributions and sparse, noisy reward signals. The rule-based nature of imperative planning [6], [5] leads to failure to adjust to novel spatial configurations. Moreover, both methods operate on single-frame input, which restricts their ability to capture holistic scene geometry and leads to poor adaptation in cluttered or unstructured settings.

In contrast, LoGoPlanner achieves stronger robustness by jointly incorporating ego-state information and multi-frame geometric reconstruction. This design ensures greater consistency in trajectory generation while providing richer spatial perception, which in turn enhances obstacle avoidance and overall navigation performance. Compared with baselines



Fig. 4. Visualization of LoGoPlanner in real-world scenarios on different robot platforms. The green curves are the planned trajectories of LoGoPlanner. Blue and grey clouds are the obstacles of the current frame and the previous frame respectively.

that benefit from external or oracle localization, LoGoPlanner improves Home SR by 27.3% points and Home SPL by 21.3% relative to ViPlanner, demonstrating that integrating implicit self-localization with geometry-aware planning yields superior closed-loop navigation.

2) *Real-world Experiments*: To evaluate the cross-platform, cross-scene, and cross-view generalization of LoGoPlanner in real-world settings, we deploy the system on three distinct robotic platforms under diverse environment configurations. For quantitative evaluation of vision-based navigation methods in real-world scenarios, we test a TurtleBot in an office scene with structured obstacles, a Unitree Go2 in a cluttered home scene containing arbitrarily shaped obstacles, and a Unitree G1 in an industrial scene with road-block obstacles, evaluating 20 trajectories. Quantitative results are shown in Table II. All algorithms run on an NVIDIA RTX 4090 GPU for cloud-based inference, with control commands transmitted to the robots in real time.

iPlanner performs relatively poorly across different embodiments because its trajectory inconsistency during obstacle avoidance often leads to collisions. In contrast, ViPlanner demonstrates better performance; however, constrained by the training data and network design of single-frame-based navigation policy, it exhibits poor performance in challenging scenarios (e.g., industrial environments with Unitree G1).

LoGoPlanner can be deployed directly without requiring visual odometry or SLAM. Despite camera jitter caused by the quadruped platform, LoGoPlanner achieves accurate self-localization and generates reliable collision-free trajectories

towards the goal. By leveraging point clouds as implicit intermediate representations, the model further reduces the sim-to-real gap, demonstrating the framework’s robust generalization, reduced deployment complexity, and readiness for direct application in scenarios with varying scene structure and camera viewpoints. For detailed demonstrations of our system, please refer to the [demo video](#).

TABLE II
REAL-WORLD RESULTS

Planner	Office	Home	Industrial
	TurtleBot	Unitree Go2	Unitree G1
	SR \uparrow	SR \uparrow	SR \uparrow
iPlanner [6]	10 (2/20)	15 (3/20)	0 (0/20)
ViPlanner [5]	50 (10/20)	45 (9/20)	0 (0/20)
LoGoPlanner	85 (17/20)	70 (14/20)	50 (10/20)

C. Ablation Study

To endow the model with self-state estimation and metric-aware perception, we introduce three auxiliary tasks in the first stage. **Odometry** supervises ego-motion estimation, **Goal** provides dynamic target updates based on self-state, and **Point Cloud** transforms historical image observations into geometric point clouds using estimated camera extrinsics. Without any auxiliary tasks, the model trained solely with end-to-end supervision exhibits basic path-planning capability. However, intermediate features lack explicit guidance, and the model may lose track of the goal when self-state estimation is inaccurate. Incorporating **Odometry**

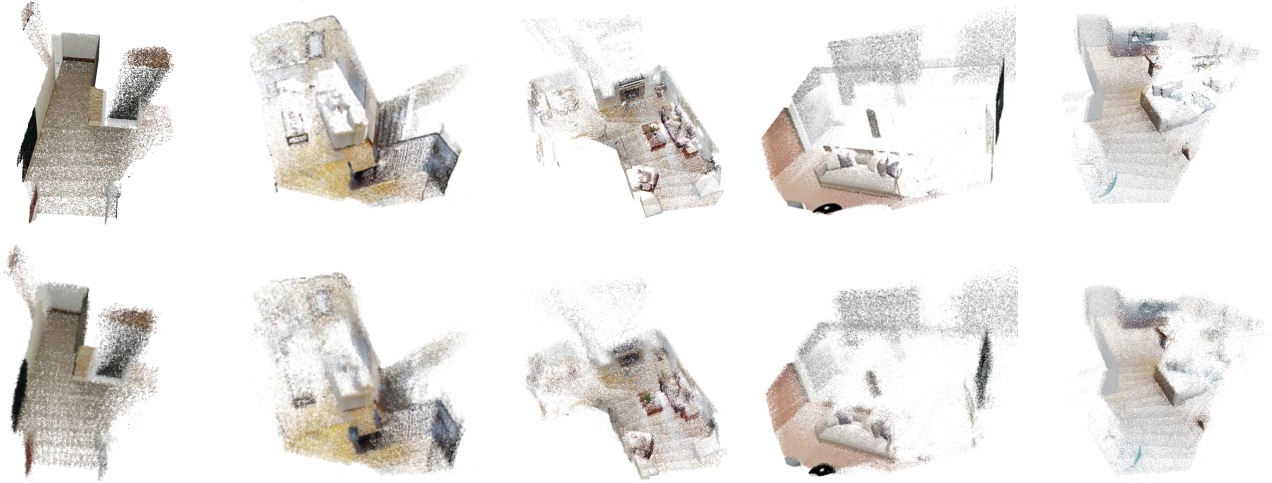


Fig. 5. Visualization of reconstruction results: the first row shows the scene point cloud of the ground truth, and the second row shows the predicted scene point cloud. The point cloud at the metric scale is predicted with the chassis of the last frame as the coordinate origin.

and **Goal** supervision improves ego-motion estimation and enhances trajectory consistency. Further adding **Point Cloud** supervision allows the model to capture the spatial relationships of obstacles beyond 2D semantics, significantly improving obstacle avoidance in trajectory generation.

TABLE III
ABLATION EXPERIMENTS ON KEY MODULES

Modules			Home		Commercial	
Odometry	Goal	Point Cloud	SR \uparrow	SPL \uparrow	SR \uparrow	SPL \uparrow
\times	\times	\times	49.5	47.0	59.4	57.0
\checkmark	\times	\times	51.3	49.7	61.2	59.2
\checkmark	\checkmark	\times	52.4	50.1	63.3	60.3
\checkmark	\checkmark	\checkmark	57.3	52.4	67.1	63.9

Our model grounds the pose estimation and reconstruction capabilities of the video geometry model into the planning policy, making the choice of geometric backbone crucial for the task. We experimented with several backbones: single-frame geometric backbones, DepthAnything [39], multi-frame geometric backbones, Video DepthAnything [21], VGGT \dagger [12] without metric scale, and a scale-injected version of VGGT [12]. During these experiments, all auxiliary tasks are retained in training.

To quantify pose estimation and planning accuracy, we define two metrics: **Navigation Error** (NE), the Euclidean distance between the robot’s final stopping position and the goal, and **Planning Error** (PE), the distance between the endpoint of the planned trajectory and the goal. Single-frame backbones provide per-frame depth perception for obstacle avoidance but lack temporal consistency. Multi-frame pretrained geometric models capture temporal correlations, yet without supervision on camera poses, they struggle to model accurate sequential pose relationships and keep planning consistency. Existing reconstruction-pretrained models offer reliable ego-motion estimation; however, after fine-tuning without depth prior, the estimated camera poses follow the correct trends but exhibit scale discrepancies.

After incorporating scale priors, our model not only achieves higher planning success rates but also demonstrates improved planning accuracy. Therefore, metric-scale supervision is still required for real-world applications.

TABLE IV
PERFORMANCE UNDER DIFFERENT VIDEO GEOMETRY BACKBONES

Backbone	Home				Commercial			
	SR \uparrow	SPL \uparrow	NE \downarrow	PE \downarrow	SR \uparrow	SPL \uparrow	NE \downarrow	PE \downarrow
DA [39]	49.9	47.1	2.51	1.48	59.9	57.4	2.49	1.49
VDA [21]	51.5	48.2	2.43	1.08	61.4	58.8	2.15	1.08
VGGT \dagger [12]	54.9	50.4	2.35	0.87	62.4	57.7	2.31	1.18
VGGT [12]	57.3	52.4	2.24	0.55	67.1	63.9	2.07	0.59

\dagger : without injecting depth as the scale prior.

VI. CONCLUSIONS

We proposed LoGoPlanner, a localization-grounded end-to-end navigation framework that unifies metric-aware pose estimation, long-horizon scene reconstruction, and feature-level policy conditioning. By integrating implicit self-state estimation with fine-grained environmental perception, our method overcomes limitations of traditional modular pipelines and prior end-to-end approaches that rely on external localization. Experiments in both simulation and real-world scenarios demonstrate that LoGoPlanner achieves superior trajectory planning and obstacle avoidance, while generalizing robustly across diverse embodiments, view-points, and environments. This work highlights the potential of grounding navigation policies in geometric and metric-aware priors, pointing toward more autonomous, reliable, and adaptable robotic navigation in unstructured real-world settings.

A key limitation of LoGoPlanner is that imitation learning on offline static trajectories limits its ability to handle highly dynamic scenes. Integrating online adaptation, such as 3D point tracking [40] or reinforcement learning-based fine-tuning [41], could improve responsiveness to moving obstacles and evolving environments.

REFERENCES

- [1] C. Cao, H. Zhu, F. Yang, Y. Xia, H. Choset, J. Oh, and J. Zhang, “Autonomous exploration development environment and the planning algorithms,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8921–8928.
- [2] D. D. Fan, K. Otsu, Y. Kubo, A. Dixit, J. Burdick, and A.-A. Agha-Mohammadi, “Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation,” *arXiv preprint arXiv:2103.02828*, 2021.
- [3] L. Wellhausen and M. Hutter, “Rough terrain navigation for legged robots using reachability planning and template learning,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6914–6921.
- [4] I. Gog, S. Kalra, P. Schafhalter, M. A. Wright, J. E. Gonzalez, and I. Stoica, “Pylot: A modular platform for exploring latency-accuracy tradeoffs in autonomous vehicles,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8806–8813.
- [5] P. Roth, J. Nubert, F. Yang, M. Mittal, and M. Hutter, “Viplanner: Visual semantic imperative learning for local navigation,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 5243–5249.
- [6] F. Yang, C. Wang, C. Cadena, and M. Hutter, “iPlanner: Imperative Path Planning,” in *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [7] D. Shah, A. Sridhar, A. Bhorkar, N. Hirose, and S. Levine, “Gnm: A general navigation model to drive any robot,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7226–7233.
- [8] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra, “Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames,” *arXiv preprint arXiv:1911.00357*, 2019.
- [9] W. Cai, J. Peng, Y. Yang, Y. Zhang, M. Wei, H. Wang, Y. Chen, T. Wang, and J. Pang, “Navdp: Learning sim-to-real navigation diffusion policy with privileged information guidance,” *arXiv preprint arXiv:2505.08712*, 2025.
- [10] A. Sridhar, D. Shah, C. Glossop, and S. Levine, “Nomad: Goal masked diffusion policies for navigation and exploration,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 63–70.
- [11] D. Shah, A. Sridhar, N. Dashora, K. Stachowicz, K. Black, N. Hirose, and S. Levine, “Vint: A foundation model for visual navigation,” in *Conference on Robot Learning*. PMLR, 2023, pp. 711–733.
- [12] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, “Vggg: Visual geometry grounded transformer,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 5294–5306.
- [13] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang *et al.*, “Planning-oriented autonomous driving,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 17 853–17 862.
- [14] J. Peng, T. Wang, J. Pang, and Y. Shen, “Towards latency-aware 3d streaming perception for autonomous driving,” *arXiv preprint arXiv:2504.19115*, 2025.
- [15] A. Loquercio, E. Kaufmann, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, “Learning high-speed flight in the wild,” *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021.
- [16] A. Sadat, S. Casas, M. Ren, X. Wu, P. Dhawan, and R. Urtasun, “Perceive, predict, and plan: Safe motion planning through interpretable semantic representations,” in *European Conference on Computer Vision*. Springer, 2020, pp. 414–430.
- [17] M. Wei, C. Wan, J. Peng, X. Yu, Y. Yang, D. Feng, W. Cai, C. Zhu, T. Wang, J. Pang *et al.*, “Ground slow, move fast: A dual-system foundation model for generalizable vision-and-language navigation,” *arXiv preprint arXiv:2512.08186*, 2025.
- [18] Z. Xu, X. Han, H. Shen, H. Jin, and K. Shimada, “Navrl: Learning safe flight in dynamic environments,” *IEEE Robotics and Automation Letters*, 2025.
- [19] Y. Wang, J. Zhou, H. Zhu, W. Chang, Y. Zhou, Z. Li, J. Chen, J. Pang, C. Shen, and T. He, “ π^3 : Scalable permutation-equivariant visual geometry learning,” *arXiv preprint arXiv:2507.13347*, 2025.
- [20] Q. Wang, Y. Zhang, A. Holynski, A. A. Efros, and A. Kanazawa, “Continuous 3d perception model with persistent state,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 10 510–10 522.
- [21] S. Chen, H. Guo, S. Zhu, F. Zhang, Z. Huang, J. Feng, and B. Kang, “Video depth anything: Consistent depth estimation for super-long videos,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 22 831–22 840.
- [22] O. Siméoni, H. V. Vo, M. Seitzer, F. Baldassarre, M. Oquab, C. Jose, V. Khalidov, M. Szafraniec, S. Yi, M. Ramamonjisoa *et al.*, “Dinov3,” *arXiv preprint arXiv:2508.10104*, 2025.
- [23] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, “Monoslam: Real-time single camera slam,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 29, no. 6, pp. 1052–1067, 2007.
- [24] G. Klein and D. Murray, “Parallel tracking and mapping on a camera phone,” in *2009 8th IEEE International Symposium on Mixed and Augmented Reality*. IEEE, 2009, pp. 83–86.
- [25] J. Engel, V. Koltun, and D. Cremers, “Direct sparse odometry,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [26] J. Engel, T. Schöps, and D. Cremers, “Lsd-slam: Large-scale direct monocular slam,” in *European conference on computer vision*. Springer, 2014, pp. 834–849.
- [27] C. Godard, O. Mac Aodha, M. Firman, and G. J. Brostow, “Digging into self-supervised monocular depth estimation,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 3828–3838.
- [28] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, “Orb-slam: A versatile and accurate monocular slam system,” *IEEE transactions on robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [29] A. Kendall, M. Grimes, and R. Cipolla, “Posenet: A convolutional network for real-time 6-dof camera relocalization,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2938–2946.
- [30] F. Wimbauer, N. Yang, L. Von Stumberg, N. Zeller, and D. Cremers, “Monorec: Semi-supervised dense reconstruction in dynamic environments from a single moving camera,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6112–6122.
- [31] Y. Wei, S. Lu, F. Han, R. Xiong, and Y. Wang, “Bev-odom: Reducing scale drift in monocular visual odometry with bev representation,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 349–356.
- [32] S. Shah, N. Rajyaguru, C. D. Singh, C. Metzler, and Y. Aloimonos, “Codedvo: Coded visual odometry,” *IEEE Robotics and Automation Letters*, 2024.
- [33] T. Qin, P. Li, and S. Shen, “Vins-mono: A robust and versatile monocular visual-inertial state estimator,” *IEEE transactions on robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [34] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby *et al.*, “Dinov2: Learning robust visual features without supervision,” *arXiv preprint arXiv:2304.07193*, 2023.
- [35] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu, “Roformer: Enhanced transformer with rotary position embedding,” *Neurocomputing*, vol. 568, p. 127063, 2024.
- [36] I.-N. D. contributors, “Interndata-n1 dataset,” <https://huggingface.co/datasets/InternRobotics/InternData-N1>, 2025, accessed: 2025-09-15.
- [37] H. Wang, J. Chen, W. Huang, Q. Ben, T. Wang, B. Mi, T. Huang, S. Zhao, Y. Chen, S. Yang, P. Cao, W. Yu, Z. Ye, J. Li, J. Long, Z. Wang, H. Wang, Y. Zhao, Z. Tu, Y. Qiao, D. Lin, and P. Jiangmiao, “Grutopia: Dream general robots in a city at scale,” in *arXiv*, 2024.
- [38] C. Campos, R. Elvira, J. J. Gómez, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM,” *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [39] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, “Depth anything v2,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 21 875–21 911, 2024.
- [40] N. Karaev, I. Rocco, B. Graham, N. Neverova, A. Vedaldi, and C. Rupprecht, “Cotracker: It is better to track together,” in *European conference on computer vision*. Springer, 2024, pp. 18–35.
- [41] A. Z. Ren, J. Lidard, L. L. Ankile, A. Simeonov, P. Agrawal, A. Majumdar, B. Burchfiel, H. Dai, and M. Simchowitz, “Diffusion policy optimization,” *arXiv preprint arXiv:2409.00588*, 2024.