

APPLV: Adaptive Planner Parameter Learning from Vision-Language-Action Model

Yuanjie Lu¹, Beichen Wang¹, Zhengqi Wu², Yang Li³, Xiaomin Lin², Chengzhi Mao³, and Xuesu Xiao¹

Abstract—Autonomous navigation in highly constrained environments remains challenging for mobile robots. Classical navigation approaches offer safety assurances but require environment-specific parameter tuning; end-to-end learning bypasses parameter tuning but struggles with precise control in constrained spaces. To this end, recent robot learning approaches automate parameter tuning while retaining classical systems’ safety, yet still face challenges in generalizing to unseen environments. Recently, Vision-Language-Action (VLA) models have shown promise by leveraging foundation models’ scene understanding capabilities, but still struggle with precise control and inference latency in navigation tasks. In this paper, we propose Adaptive Planner Parameter Learning from Vision-Language-Action Model (APPLV). Unlike traditional VLA models that directly output actions, APPLV leverages pre-trained vision-language models with a regression head to predict planner parameters that configure classical planners. We develop two training strategies: supervised learning fine-tuning from collected navigation trajectories and reinforcement learning fine-tuning to further optimize navigation performance. We evaluate APPLV across multiple motion planners on the simulated Benchmark Autonomous Robot Navigation (BARN) dataset and in physical robot experiments. Results demonstrate that APPLV outperforms existing methods in both navigation performance and generalization to unseen environments.

I. INTRODUCTION

Autonomous mobile robot navigation in complex, highly constrained environments remains a fundamental challenge in robotics, with critical applications spanning warehouse logistics, autonomous delivery, and service robotics. The central problem is enabling robots to safely and efficiently navigate narrow passages and cluttered spaces with extremely low clearance while maintaining robust performance across diverse, previously unseen scenarios.

Traditional navigation systems have established themselves as the industry standard due to their reliability and safety assurances. However, they suffer from a critical weakness: performance heavily depends on parameter configurations that must be precisely tuned for each deployment scenario, including velocity limits, cost weights, sampling densities, and safety margins. This creates practical barriers: it requires expertise most end users lack, is time-consuming and environment-specific, and static parameters cannot adapt to varying conditions within a single environment. To address parameter sensitivity, end-to-end learning methods directly

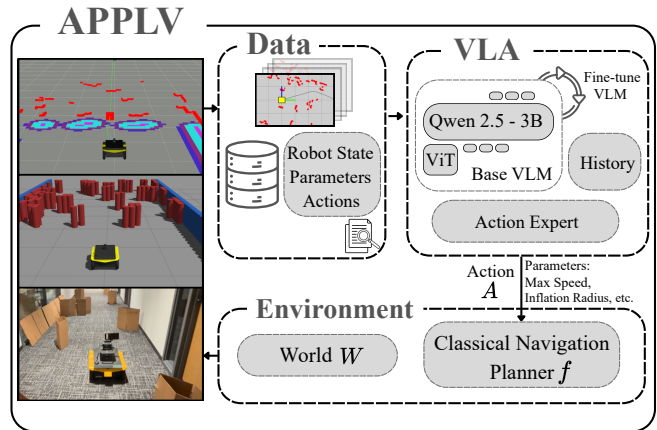


Fig. 1: APPLV instantiates the APPL paradigm with a VLA model to dynamically adjust a classical navigation planner’s parameters, learned from data collected in simulation.

map sensory inputs to control commands, bypassing parameter tuning entirely. While conceptually elegant, these approaches sacrifice the interpretability and safety assurances of classical systems and struggle to generalize beyond the training distribution. In highly constrained environments where tight clearances demand precision, real-world sensor noise and actuation uncertainties make reliable generalization even more challenging.

Recognizing the limitations of both extremes, hybrid approaches in robot learning have emerged to combine the safety of classical navigation with learning-based adaptability. The Adaptive Planner Parameter Learning (APPL) family [1] exemplifies this paradigm by automating online parameter selection, demonstrating improvements over both manual tuning and end-to-end learning. However, these methods struggle to generalize to unseen environments, with performance degrading even in qualitatively similar scenarios. Moreover, requiring human interactions [2]–[4] or random exploration [5], their overall navigation performance still leaves considerable room for improvement. In challenging spaces like narrow corridors or dense clutter, they often exhibit suboptimal behavior, either showing excessive caution that causes slow traversal or making aggressive choices that compromise safety.

Recently, vision-language foundation models have demonstrated remarkable capabilities in understanding complex visual scenes and reasoning about spatial relationships. Building on these advances, Vision-Language-Action (VLA) models have emerged as a promising paradigm for robot learning,

¹Department of Computer Science, George Mason University, Virginia, USA

²Department of Engineering Science, University of South Florida, Florida, USA

³Department of Computer Science, Rutgers University, New Jersey, USA

leveraging pre-trained Vision-Language Models (VLMs) to directly map visual observations to robot actions. However, VLA models face challenges when applied to constrained navigation: they struggle with the centimeter-level precision required in tight spaces and suffer from high inference latency unsuitable for real-time control.

Motivated by these limitations, we propose Adaptive Planner Parameter Learning from Vision-Language-Action Model (APPLV, Fig. 1), which leverages a vision-language-action model to predict navigation planner parameters rather than actions directly. Specifically, a VLM processes custom observations to extract multi-layer hidden states, while a history encoder captures temporal context from prior frames. These representations are fed into an action expert comprising a Dense Prediction Transformer (DPT)-style regression head [6] that fuses multi-layer features and temporal context to predict parameters such as velocity limits, sampling densities, and cost weights. These parameters configure a classical navigation planner to generate executable robot actions, enabling fine-grained environmental adaptation while maintaining safety assurances and computational efficiency. We propose two training strategies for APPLV: supervised learning fine-tuning from collected navigation trajectories (APPLV-SL) and reinforcement learning fine-tuning (APPLV-RLFT) to further optimize navigation performance. We evaluate APPLV across four local planners on the simulated Benchmark Autonomous Robot Navigation (BARN) dataset [7], [8] and in physical robot experiments, demonstrating superior performance over existing parameter learning methods in both navigation performance and generalization to unseen environments.

II. RELATED WORK

We review related work on autonomous robot navigation, tracing the evolution from classical parameter tuning to end-to-end learning methods, and more recently to hybrid approaches and foundation models.

Classical navigation systems require manual parameter tuning for each deployment [9], [10]. While automated optimization approaches using fuzzy logic [11] and neuro-fuzzy systems [12] reduce manual effort, they remain fundamentally limited due to their assumption that a single set of optimal parameters should fit everywhere in a deployment environment. This has led researchers to explore end-to-end learning methods that bypass parameter tuning entirely.

To bypass parameter tuning, end-to-end navigation methods directly map sensory inputs to control commands. Imitation learning approaches [13], [14] train navigation policies from expert demonstrations, and reinforcement learning methods [15]–[18] learn through trial-and-error interaction with environments. However, these approaches sacrifice the interpretability and safety assurances of classical systems. Moreover, they struggle to generalize beyond training distributions, particularly in highly constrained environments where tight clearances demand centimeter-level precision and real-world sensor noise creates additional challenges.

To address the limitations of both parameter tuning and end-to-end learning, hybrid approaches combine classical navigation methods with learning-based parameter selection. Neurofuzzy methods [19] use neural networks to tune controller parameters. Reinforcement learning approaches adapt cost functions [20] or dynamics parameters [21] to varying environmental complexity. The APPL family [1] enables dynamic parameter selection during navigation. APPLD [2] learns parameter sets from demonstrations, APPLI [3] from corrective interventions, APPLE [4] from evaluative feedback, and APPLR [5] uses reinforcement learning to learn a parameter policy. Although these approaches improve upon both manual parameter tuning and end-to-end learning, significant challenges remain: they struggle to generalize to unseen environments, and overall navigation performance in highly constrained spaces still leaves substantial room for improvement [8], [22]–[24].

Recently, VLMs trained on large-scale image-text data have demonstrated strong visual understanding capabilities. Building on this success, VLA models [25]–[27] adapt these foundation models to robotics by fine-tuning them to directly predict robot actions. For navigation, approaches such as LM-Nav [28], NavILA [29], QUAR-VLA [30], and OmniVLA [31] apply VLMs to instruction following and multi-modal goal specifications. While these foundation model approaches demonstrate strong language grounding and scene understanding capabilities, their application to constrained navigation has not been investigated before and suffers from substantial limitations. Primarily designed and evaluated in open spaces with high computational tolerance, they struggle to achieve the centimeter-level precision required in narrow passages and cluttered environments. Moreover, their large model sizes result in high inference latency, causing control delays unsuitable for real-time navigation.

Our approach, APPLV, explores an alternative direction: instead of directly predicting actions, we leverage vision-language models to predict low-level navigation parameters for classical planners, enabling navigation in highly constrained spaces. By operating at the parameter level, APPLV provides fine-grained adaptation to environmental constraints while preserving the safety assurances and interpretability of classical systems. Unlike direct action prediction that demands high-frequency inference at every control step, parameter prediction operates at a much lower frequency since classical planners can autonomously produce real-time actions under the same set of parameters, significantly relaxing real-time computational constraints.

III. PRELIMINARIES

APPLV leverages pre-trained VLMs with superior reasoning capabilities further fine-tuned on robot navigation data within the APPL paradigm to enhance the APPL parameter policy’s ability to dynamically adjust planners’ parameters.

A. The APPL Paradigm

The APPL paradigm assumes a mobile robot equipped with a classical navigation planner $f : \mathcal{X} \times \Phi \rightarrow \mathcal{A}$, where \mathcal{X} is the

planner’s state space (e.g., robot odometry, sensory inputs, and navigation goal), Φ is the parameter space (e.g., velocity limits, cost coefficients, and inflation radius), and \mathcal{A} is the action space (e.g., linear velocity v and angular velocity ω). Given a state $x \in \mathcal{X}$ and planner parameters $\phi \in \Phi \subset \mathbb{R}^d$ with d depending on the target planner, the planner generates an action $a = f(x; \phi) = f_\phi(x)$. APPL formulates parameter learning in a meta-MDP with a meta-environment \mathcal{E} that combines the physical world \mathcal{W} and planner f . The meta-state $s_t = (x_t, \phi_{t-1}) \in \mathcal{S} = \mathcal{X} \times \Phi$ includes both the planner’s state and previous parameters, while the meta-action is the parameter configuration $\phi_t \in \Phi$. An APPL agent learns a parameter policy $\pi_p : \mathcal{S} \rightarrow \Phi$ that selects parameters to maximize expected cumulative reward, instead of the conventional motion policy $\pi_m : \mathcal{X} \rightarrow \mathcal{A}$ that issues raw motor commands. Building upon the classical navigation planner f , APPL inherits f ’s safety and explainability.

B. Vision-Language Models

VLMs learn joint representations of visual and linguistic information by pre-training on large-scale image-text datasets. A typical VLM consists of a visual encoder, a projection module, and a Large Language Model (LLM) backbone, trained end-to-end with a next-token prediction objective. Qwen2.5-VL [32] follows this architecture, encoding image inputs as patch embeddings that are projected and concatenated with text tokens before being processed by the LLM. VLMs can be efficiently adapted to downstream tasks via Low-Rank Adaptation (LoRA) [33], which inserts trainable low-rank matrices into the transformer layers while keeping pre-trained weights frozen.

IV. APPLV

Fig. 2 illustrates the APPLV architecture. The meta-state is $s_t = (x_t, \phi_{t-1})$, where the planner’s state x_t comprises a current custom image I_t , historical images from previous frames $\{I_{t-k}\}_{k=1}^K$, and current robot state including linear and angular velocities, while ϕ_{t-1} are the previous planner parameters. A Qwen2.5-VL-3B model processes the current image with a language prompt, extracting multi-layer hidden states. These features are concatenated with features from the history encoder and fed into a DPT regression head to predict planner parameters ϕ_t , which configure the classical navigation planner f to generate motion control commands.

A. The APPLV Model Architecture

1) *Training Dataset*: Each training sample consists of a custom image I_t , historical frames $\{I_{t-k}\}_{k=1}^K$ for temporal context, robot current linear and angular velocity state encoded in the text prompt, and corresponding planner parameters ϕ_t . The custom image (Fig. 2, top left) is a top-down RGB representation of the robot’s local environment with:

- **Gray background**: environment background in gray;
- **Red overlay**: laser scan of obstacles;
- **Blue path**: global path from the robot to its goal; and
- **Robot marker**: robot footprint in the base link frame.

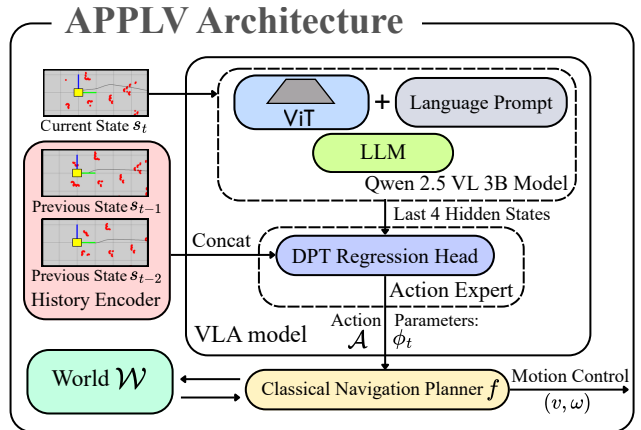


Fig. 2: APPLV Architecture. A VLA model processes current and historical states to predict planner parameters, which-configure a classical navigation planner to generate motion control commands.

Samples are stored in JSON format and shuffled during training. The demonstration dataset for supervised learning comprises data from two sources: (1) hand-crafted heuristic rules designed by an expert roboticist that map perceived environmental conditions (e.g., obstacle density and corridor width) to predefined parameter configurations, and (2) policies learned by the APPLR [5] baseline method. Both data collection approaches are detailed in Section V-A.

2) *Vision-Language Backbone*: Qwen2.5-VL-3B is employed to extract visual-semantic representations from custom images and text prompts. The model consists of a Vision Transformer (ViT) encoder [34] and a language model backbone. Unlike standard VLM applications that generate text outputs, intermediate hidden states are extracted for parameter regression. Hidden states from the last four transformer layers are extracted as $\{\mathbf{h}_t^{(l)}\}_{l=L-3}^L$, where each $\mathbf{h}_t^{(l)} \in \mathbb{R}^{N \times 256}$ represents N tokens. These multi-layer features capture spatial patterns at different levels of abstraction. During training, the vision encoder is frozen, while the language model is adapted using LoRA with rank $r = 64$ and scaling factor $\alpha = 128$ to enable parameter-efficient fine-tuning.

3) *History Encoder*: To incorporate temporal context, a lightweight history encoder processes prior visual frames. A convolutional network extracts spatial features from each frame independently. The frame-level features are augmented with learnable positional embeddings and fed into a temporal transformer encoder that models sequential dependencies across frames. The encoded historical representation is fused with the current frame’s VLM features in the DPT head, enabling temporally-aware parameter predictions.

4) *DPT Regression Head*: The DPT regression head [6] serves as the action expert, fusing multi-layer hidden states with encoded historical features to predict planner parameters. The architecture reassembles tokens from different transformer stages into feature representations and progressively combines them through convolutional fusion. Each layer’s hidden states $\mathbf{h}_t^{(l)}$ are projected to a unified feature

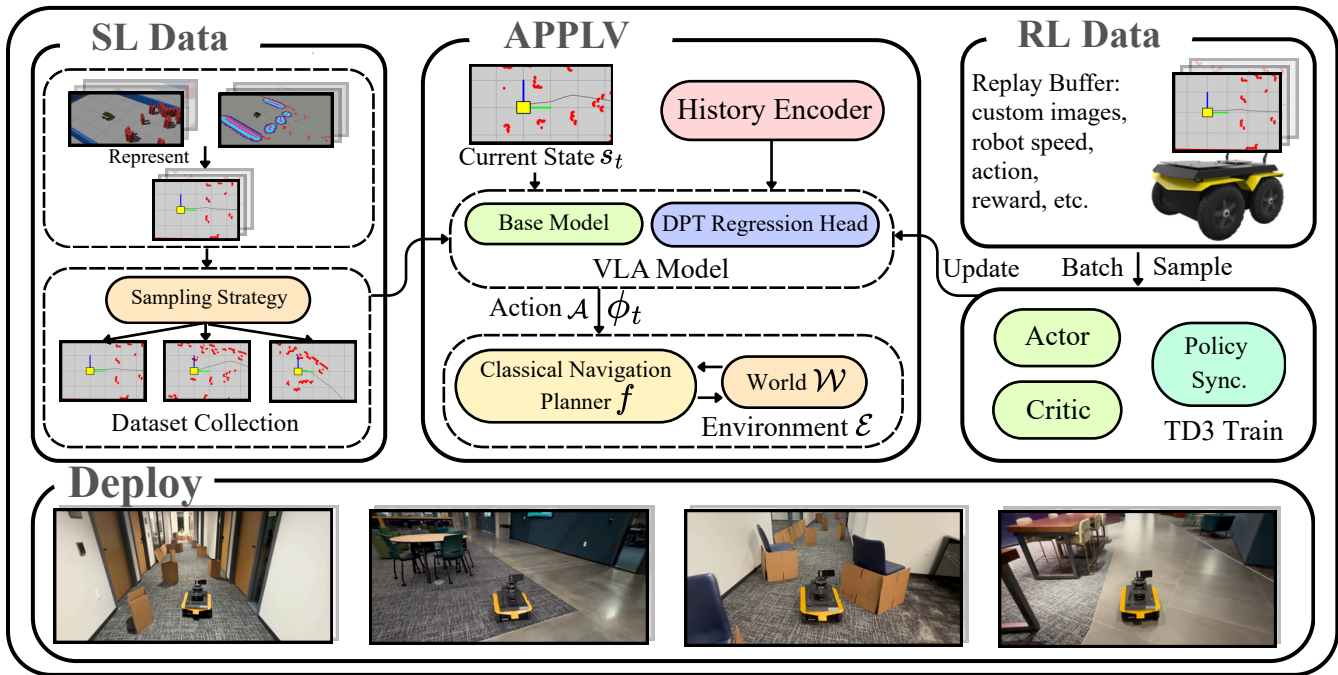


Fig. 3: APPLV Training and Deployment Pipeline. Left: Supervised learning data collection with representation and sampling. Middle: APPLV architecture with VLA model and classical planner. Right: Reinforcement Learning fine-tuning with TD3. Bottom: Physical deployment in natural cluttered environments.

space: $\mathbf{f}_t^{(l)} \in \mathbb{R}^{N \times 256}$. Features are fused progressively from deeper to shallower layers using residual convolutional blocks, with the historical context concatenated during fusion. The fused representation undergoes attention-weighted pooling to produce a global feature vector $\mathbf{z}_t \in \mathbb{R}^{256}$. A final MLP regresses the parameter vector: $\phi_t = \text{MLP}_{\text{reg}}(\mathbf{z}_t)$, where $\phi_t \in \mathbb{R}^d$ with d depending on the target planner.

5) *Classical Navigation Planner*: The classical navigation planner receives predicted parameters ϕ_t and generates motion control commands (v_t, ω_t) . Different planner implementations share a common parameterized structure where ϕ_t configures behavior-critical aspects such as velocity limits, cost function weights, planning horizons, and inflation radius. Based on these parameters, the planner selects optimal actions while ensuring kinodynamic feasibility and collision avoidance. This design allows the VLA model to adapt planning behavior across diverse scenarios and planner types through learned parameter selection.

B. Training Strategies

1) *Supervised Learning (APPLV-SL)*: The base VLA model is trained via Behavior Cloning (BC) on demonstration trajectories (Fig. 3 left). As illustrated, raw demonstration data is first processed through a representation step that converts environmental observations into custom images. A sampling strategy then selects diverse training samples to construct the final dataset. Given a dataset \mathcal{D} of such samples, the training objective minimizes mean squared error:

$$\mathcal{L}_{\text{SL}} = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \|\phi_i - \phi_i^*\|_2^2,$$

where ϕ_i is the predicted parameter for sample i and ϕ_i^* is the corresponding ground-truth parameter from demonstration data.

2) *Reinforcement Learning Fine-Tuning (APPLV-RLFT)*: The second training strategy employs off-policy reinforcement learning (Fig. 3 right). The VLA model can be initialized from the supervised pretrained weights or trained from scratch; we adopt the former for accelerated convergence. The input and output spaces remain identical to supervised learning. A distributed collection framework deploys multiple parallel collectors, each running the latest policy to interact with an independent simulation environment. Experiences are aggregated into a shared replay buffer. The reward function encourages effective navigation through four components: (1) progress reward r_{progress} measuring distance reduction toward the goal, (2) collision penalty $r_{\text{collision}}$ enforcing safety constraints, (3) time penalty r_{time} promoting efficient task completion, and (4) obstacle avoidance reward r_{obstacle} encouraging safe proximity to environmental obstacles. The total reward is: $r_t = r_{\text{progress}} + r_{\text{collision}} + r_{\text{time}} + r_{\text{obstacle}}$. The VLA model serves as the actor, predicting planner parameters from visual and state inputs. A separate critic network estimates the value of the predicted parameters. Training follows the Twin Delayed Deep Deterministic Policy Gradient (TD3) algorithm [35]: the critic is updated every step by minimizing temporal difference error between predicted and target Q-values, while the actor is updated every two steps by maximizing the critic’s Q-value estimate. Periodic policy synchronization ensures collectors use recent policies, balancing exploration and exploitation.



Fig. 4: Physical deployment of two BARN Challenge environments

V. EXPERIMENTS

The BARN Challenge benchmark focuses on highly constrained scenarios with narrow passages and dense obstacles. We integrate APPLV with four representative local planners: DWA [36], TEB [37], MPPI [38], and DDP [39]. APPLV is compared against baselines spanning heuristics, BC, reinforcement learning, and zero-shot foundation model approaches. Experiments are conducted across 300 simulated BARN environments [7], with additional real-world validation on a Clearpath Jackal robot.

A. Experimental Setup

1) *Navigation Planners*: To demonstrate that APPLV generalizes across different navigation planners, we integrate it with four representative local planners: DWA [36], which samples velocity commands within the robot’s dynamic window; TEB [37], which optimizes an elastic band of timed poses under kinodynamic constraints; MPPI [38], a sampling-based planner that computes optimal controls via weighted averaging of stochastic rollouts; and DDP [39], a recently proposed dynamics-based planner that achieved second place in the 2025 BARN Challenge [8]. DWA and TEB are integrated from the ROS `move_base` stack, while MPPI and DDP are custom C++ implementations.

2) *Baseline Methods*: We compare APPLV against four baseline methods for navigation parameter prediction:

Heuristic Expert represents the traditional approach to parameter configuration in classical navigation systems. An

expert roboticist designs hand-crafted rules that map perceived environmental conditions, such as obstacle density and corridor width, to predefined parameter configurations. This baseline reflects common practice in real-world deployments and establishes a practical comparison for learning-based methods.

APPLR [5], a representative method from the APPL family [1], trains a TD3-based reinforcement learning policy that maps laser scan observations to navigation parameters. This purely RL-based approach has demonstrated substantial improvements over static parameter configurations in classical planners like DWA, serving as a strong baseline for learned parameter adaptation.

Transformer BC serves as a controlled baseline to isolate the contribution of vision-language pre-training. We train a transformer model from scratch with a comparable number of parameters and the same training data as APPLV, using supervised learning on collected demonstrations to predict navigation parameters. Unlike APPLV which fine-tunes a pre-trained VLM, Transformer BC uses random initialization, allowing us to quantify the performance gain from leveraging vision-language pre-training.

Zero-Shot VLM evaluates whether foundation models can solve navigation parameter prediction without task-specific training. We prompt GPT-4o [40] with custom images and robot state to directly generate navigation parameters as text output, relying entirely on the model’s zero-shot spatial reasoning and numerical estimation capabilities. This baseline demonstrates the performance gap between zero-shot prompting and task-specific fine-tuning, validating the necessity of adapting foundation models to navigation tasks.

3) *Experimental Environments*: Our evaluation encompasses both simulated and real-world scenarios. The BARN benchmark [7] provides 300 highly constrained environments with diverse obstacle configurations generated via cellular automata for training (Fig. 3, middle left). An additional set of 300 entirely different environments is used for testing to evaluate generalization to unseen scenarios. Physical experiments are conducted in indoor test courses using a Clearpath Jackal robot (Fig. 3 bottom; Fig. 4).

4) *Pre-Train Dataset Collection*: Training data is collected using both Heuristic Expert and APPLR parameters in simulation environments. For each of the four planners, we execute 1000 trajectory attempts per environment across all 300 BARN training environments, collecting both successful and failed navigation experiences. During each trajectory execution, we record per-step observations at 0.5-second intervals, including custom images, robot velocities, and corresponding planner parameters. Each trajectory contains multiple time steps, forming a sequence of navigation experiences. To ensure dataset quality and balance, we employ two sampling strategies: difficulty-aware sampling and motion-aware filtering. Difficulty-aware sampling classifies environments into four difficulty tiers based on average navigation score from the BARN Challenge. We apply progressively relaxed score thresholds across tiers: easier environments retain only high-scoring trajectories to focus on successful behaviors,

TABLE I: Performance Comparison of Methods Grouped by Planners on 300 Test BARN Environments.

Planner	Method	Success (%) \uparrow	Avg. Time (s) \downarrow	Avg. Score \uparrow	Collision (%) \downarrow	Timeout (%) \downarrow
DWA	APPLR	73.15	27.38	0.296	11.87	14.98
	Heuristic Expert	82.47	25.83	0.328	07.50	10.03
	Transformer BC	83.03	27.58	0.323	08.21	08.75
	Zero-Shot VLM	81.00	31.27	0.301	10.00	09.00
	APPLV-SL	86.10	21.58	0.357	07.51	06.39
	APPLV-RLFT	87.20	18.68	0.374	06.51	06.29
MPPI	APPLR	78.53	24.73	0.356	07.40	14.07
	Heuristic Expert	84.48	23.41	0.365	14.17	01.35
	Transformer BC	83.68	20.42	0.378	06.52	09.80
	Zero-Shot VLM	85.24	22.02	0.367	08.02	06.74
	APPLV-SL	88.93	17.82	0.415	09.84	01.22
	APPLV-RLFT	89.70	16.75	0.434	07.30	03.00
TEB	APPLR	79.17	20.29	0.362	05.53	15.30
	Heuristic Expert	86.61	19.11	0.388	00.58	12.81
	Transformer BC	90.25	20.35	0.383	02.10	07.64
	Zero-Shot VLM	89.31	16.64	0.398	01.07	09.61
	APPLV-SL	90.00	17.61	0.429	00.78	09.22
	APPLV-RLFT	90.30	12.51	0.441	08.90	00.70
DDP	APPLR	85.35	15.66	0.404	07.93	06.72
	Heuristic Expert	89.50	16.09	0.418	02.38	08.12
	Transformer BC	85.57	18.16	0.411	02.36	12.05
	Zero-Shot VLM	92.50	16.18	0.417	00.50	04.00
	APPLV-SL	92.68	15.28	0.440	01.66	05.65
	APPLV-RLFT	94.34	13.63	0.463	01.98	03.68

Each environment is tested 3 times. Bold indicates the best result within each planner group.

while harder environments gradually lower the threshold to admit broader trajectory distributions, ensuring sufficient training data diversity in challenging scenarios where high-scoring trajectories are scarce. Motion-aware filtering evaluates whether each frame makes progress toward the goal by measuring the distance reduction to the target. Frames with forward progress (i.e., distance to goal decreases) are fully retained as they demonstrate effective navigation behavior. In contrast, stagnant or rotational frames (i.e., distance to goal unchanged or increasing) are subsampled at approximately 10% retention rate, as they provide limited learning value while still capturing obstacle avoidance behaviors. After applying both sampling strategies, we obtain approximately 30,000 high-quality training samples per planner.

5) *Training Parameters Configuration*: The APPLV model processes 600×400 pixel custom images, which together with text prompts occupy approximately 601 tokens as input to the VLM. The model consists of three main components: LoRA adapters (330M parameters) applied to the pre-trained vision-language backbone, a DPT-style regression head (3.8M parameters), and a history encoder (1.7M parameters). During supervised pre-training, we train the model with a batch size of 4, optimizing the LoRA adapters, history encoder, and regression head while keeping the VLM backbone frozen. For Reinforcement Learning Fine-Tuning (RLFT), we adopt a TD3-based framework where the actor network is initialized from the supervised learning checkpoint. The critic network has the same architecture and parameter count as the actor, but with independent parameters. Unlike APPLR [5], which uses a 2-layer MLP to map 721-dimensional observations (720 laser rays + goal angle) to navigation parameters, our actor processes custom images through the full APPLV

architecture with vision-language representations. We use ZeRO-2 optimization with a batch size of 42 during RLFT.

6) *Software and Hardware Configuration*: All simulation experiments are conducted on a SLURM cluster using Singularity containers [41]. The simulation environment runs Ubuntu 20.04 with ROS Noetic and Python 3.8, executing 250-300 parallel Gazebo instances. Simulations communicate with FastAPI inference servers on A100 GPU nodes. Physical experiments use a Clearpath Jackal robot with a Hokuyo LiDAR sensor (720-dimensional scans, 270° field of view). The robot’s onboard computer sends observations to an RTX 5070 Ti GPU server via FastAPI for real-time parameter prediction. Model training and inference use Python 3.10. The APPLV model is built on Qwen2.5-VL-3B. Inference latency per prediction is approximately 0.41s on RTX 5070 Ti, 0.47s on RTX 3080, and 0.27s on RTX 5090.

B. Experimental Results

We evaluate navigation performance in simulation using the BARN Challenge scoring metric [8], which balances success rate and traversal time efficiency. Physical experiments use success rate, completion progress, and traversal time as metrics.

1) *Simulated BARN Environments*: Table I presents the performance of APPLV against baseline methods across four local planners on 300 unseen BARN test environments. Both APPLV-SL and APPLV-RLFT demonstrate substantial improvements over all baselines in terms of success rate, navigation time, and overall navigation score. To isolate the contribution of vision-language pre-training, we compare APPLV-SL against Transformer BC, which uses a comparable model size and the same training data but is trained from scratch.

TABLE II: Performance in Test Physical Environments.

Planner	Method	Success \uparrow	Progress (%) \uparrow	Time (s) \downarrow
DWA	APPLR	0/6	13.00 \pm 3.2	120.00 \pm 0.0
	Heuristic Expert	0/6	37.00 \pm 4.5	120.00 \pm 0.0
	Transformer BC	0/6	40.00 \pm 3.8	120.00 \pm 0.0
	APPLV-SL	0/6	64.00 \pm 2.7	120.00 \pm 0.0
	APPLV-RLFT	1/6	75.00 \pm 3.5	95.00 \pm 2.9
MPPI	APPLR	1/6	32.00 \pm 4.1	102.00 \pm 3.6
	Heuristic Expert	3/6	73.00 \pm 2.8	80.00 \pm 4.7
	Transformer BC	2/6	71.00 \pm 3.4	61.00 \pm 2.5
	APPLV-SL	6/6	100.00 \pm 0.0	40.00 \pm 3.9
	APPLV-RLFT	6/6	100.00 \pm 0.0	34.00 \pm 2.3
TEB	APPLR	0/6	18.00 \pm 4.2	120.00 \pm 0.0
	Heuristic Expert	1/6	32.00 \pm 3.6	105.00 \pm 4.8
	Transformer BC	1/6	34.00 \pm 2.9	103.00 \pm 3.7
	APPLV-SL	2/6	64.00 \pm 4.4	83.00 \pm 2.6
	APPLV-RLFT	2/6	87.00 \pm 3.1	67.00 \pm 4.5
DDP	APPLR	0/6	22.00 \pm 3.8	120.00 \pm 0.0
	Heuristic Expert	2/6	78.00 \pm 2.4	56.00 \pm 3.3
	Transformer BC	3/6	76.00 \pm 4.6	64.00 \pm 2.7
	APPLV-SL	6/6	100.00 \pm 0.0	34.00 \pm 4.1
	APPLV-RLFT	6/6	100.00 \pm 0.0	32.00 \pm 3.5

Each environment is tested 3 times. Bold indicates the best result within each planner group.

APPLV-SL consistently outperforms Transformer BC across all planners, validating that leveraging pre-trained VLMs significantly improves performance. Comparing APPLV-SL with Zero-Shot VLM (GPT-4o prompted to output parameters) evaluates whether fine-tuning is necessary. While Zero-Shot VLM demonstrates reasonable performance, APPLV-SL achieves superior results, particularly with simpler planners. This indicates that task-specific fine-tuning remains essential despite foundation models’ spatial reasoning capabilities. APPLV-RLFT consistently improves upon APPLV-SL across all planners, demonstrating that RLFT effectively optimizes for task objectives beyond supervised imitation. Compared to the reinforcement-learning-only baseline APPLR, APPLV-RLFT achieves substantially higher success rates and faster navigation across all planners, validating that vision-language representations significantly outperform laser-scan-based approaches.

2) *Physical Experiments:* We conduct physical robot experiments in two distinct real-world environments, with each configuration tested three times. Performance is evaluated by success rate, average progress percentage, and average completion time. Progress percentage measures the distance covered toward the goal, reaching 100% upon successful arrival. Failed or timed-out trials are assigned a maximum time of 120 seconds. We exclude Zero-Shot VLM baselines from physical experiments due to their prohibitively long inference time. Table II presents the physical experiment results across four planners. APPLV consistently outperforms all three baselines (APPLR, Heuristic Expert, and Transformer BC) across all planners, with APPLV-RLFT surpassing APPLV-SL in all cases. Notably, DWA and TEB from ROS move_base exhibit significantly degraded performance in physical deployment compared to simulation, while custom C++ implementations




Fig. 5: Effect of Training Data Size (APPLV-SL).

of MPPI and DDP maintain strong performance. This discrepancy arises from fundamental architectural differences: DWA and TEB rely heavily on costmap-based global path planning, where even minor localization errors can cause severe costmap distortions. When the inflation radius becomes misconfigured due to these errors, the global planner may fail to generate valid paths, causing navigation failure. In contrast, custom MPPI and DDP implementations depend primarily on direct laser scan observations and maintain a history of global paths, making them more robust to transient localization errors in physical environments.

C. Effect of Training Data Size

We investigate the impact of training data size on APPLV-SL by varying samples from 1,500 to 5,500. Fig. 5 shows that performance does not monotonically improve with data size. All planners reach peak performance at moderate data sizes and plateau or degrade thereafter. This suggests that additional data primarily helps the VLM learn feature representations rather than memorize examples. Once core features are captured, further data provides diminishing returns.

VI. CONCLUSION

We present APPLV, which leverages vision-language models for adaptive navigation planner parameter learning. This work demonstrates that pre-trained VLMs possess powerful visual reasoning capabilities for robot navigation. By predicting planner parameters rather than direct actions, APPLV enables dynamic adaptation while maintaining safety assurances of classical planners.

Experiments across four planners on the BARN benchmark and physical robots show that APPLV consistently outperforms existing methods, including APPLR, Heuristic Expert, and Transformer BC. This validates that vision-language models can effectively understand navigation scenarios and make informed parameter decisions, offering a promising direction for adaptive robot navigation that combines foundation model strengths with classical planning reliability.

REFERENCES

- [1] X. Xiao, Z. Wang, Z. Xu, B. Liu, G. Warnell, G. Dhamankar, A. Nair, and P. Stone, “Appl: Adaptive planner parameter learning,” *Robotics and Autonomous Systems*, vol. 154, p. 104132, 2022.

- [2] X. Xiao, B. Liu, G. Warnell, J. Fink, and P. Stone, "Appld: Adaptive planner parameter learning from demonstration," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4541–4547, 2020.
- [3] Z. Wang, X. Xiao, B. Liu, G. Warnell, and P. Stone, "Appli: Adaptive planner parameter learning from interventions," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6079–6085.
- [4] Z. Wang, X. Xiao, G. Warnell, and P. Stone, "Apple: Adaptive planner parameter learning from evaluative feedback," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7744–7749, 2021.
- [5] Z. Xu, G. Dhamankar, A. Nair, X. Xiao, G. Warnell, B. Liu, Z. Wang, and P. Stone, "Applr: Adaptive planner parameter learning from reinforcement," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 6086–6092.
- [6] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 12 179–12 188.
- [7] D. Perille, A. Truong, X. Xiao, and P. Stone, "Benchmarking metric ground navigation," in *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2020, pp. 116–121.
- [8] X. Xiao, Z. Xu, S. A. Ghani, A. Datar, D. Song, P. Stone, A. Mazen, K. Yazdipaz, I. Mateyaunga, M. Faied *et al.*, "Autonomous ground navigation in highly constrained spaces: Lessons learned from the fourth barn challenge at icra 2025."
- [9] R. L. Guimarães, A. S. de Oliveira, J. A. Fabro, T. Becker, and V. A. Brenner, "Ros navigation: Concepts and tutorial," in *Robot Operating System (ROS) The Complete Reference (Volume 1)*. Springer, 2016, pp. 121–160.
- [10] K. Zheng, "Ros navigation tuning guide," in *Robot Operating System (ROS) The Complete Reference (Volume 6)*. Springer, 2021, pp. 197–226.
- [11] O. Castillo, H. Neyoy, J. Soria, P. Melin, and F. Valdez, "A new approach for dynamic fuzzy logic parameter tuning in ant colony optimization and its application in fuzzy control of a mobile robot," *Applied soft computing*, vol. 28, pp. 150–159, 2015.
- [12] D. Teso-Fz-Betoño, E. Zulueta, U. Fernandez-Gamiz, A. Saenz-Aguirre, and R. Martinez, "Predictive dynamic window approach development with artificial neural fuzzy inference improvement," *Electronics*, vol. 8, no. 9, p. 935, 2019.
- [13] B. César-Tondreau, G. Warnell, E. Stump, K. Kochersberger, and N. R. Waytowich, "Improving autonomous robotic navigation using imitation learning," *Frontiers in Robotics and AI*, vol. 8, p. 627730, 2021.
- [14] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4693–4700.
- [15] T. Xu, C. Pan, and X. Xiao, "Reinforcement learning for wheeled mobility on vertically challenging terrain," in *2024 IEEE International Symposium on Safety Security Rescue Robotics (SSRR)*. IEEE, 2024, pp. 125–130.
- [16] A. Faust, K. Oslund, O. Ramirez, A. Francis, L. Tapia, M. Fiser, and J. Davidson, "Prm-rl: Long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 5113–5120.
- [17] H. Shi, L. Shi, M. Xu, and K.-S. Hwang, "End-to-end navigation strategy with deep reinforcement learning for mobile robots," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 4, pp. 2393–2402, 2019.
- [18] L. Wang, T. Xu, Y. Lu, and X. Xiao, "Reward training wheels: Adaptive auxiliary rewards for robotics reinforcement learning," *arXiv preprint arXiv:2503.15724*, 2025.
- [19] A. Zhu and S. X. Yang, "Neurofuzzy-based approach to mobile robot navigation in unknown environments," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 4, pp. 610–621, 2007.
- [20] X. Xiao, T. Zhang, K. M. Choromanski, T.-W. E. Lee, A. Francis, J. Varley, S. Tu, S. Singh, P. Xu, F. Xia, S. M. Persson, D. Kalashnikov, L. Takayama, R. Frostig, J. Tan, C. Parada, and V. Sindhwani, "Learning model predictive controllers with real-time attention for real-world navigation," in *Proceedings of The 6th Conference on Robot Learning (CoRL)*, ser. Proceedings of Machine Learning Research, vol. 205, 2023, pp. 1708–1721.
- [21] Y. Lu, M. Mao, T. Xu, L. Wang, X. Lin, and X. Xiao, "Adaptive dynamics planning for robot navigation," in *2026 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2026.
- [22] X. Xiao, Z. Xu, Z. Wang, Y. Song, G. Warnell, P. Stone, T. Zhang, S. Ravi, G. Wang, H. Karnan *et al.*, "Autonomous ground navigation in highly constrained spaces: Lessons learned from the benchmark autonomous robot navigation challenge at icra 2022 [competitions]," *IEEE Robotics & Automation Magazine*, vol. 29, no. 4, pp. 148–156, 2022.
- [23] X. Xiao, Z. Xu, G. Warnell, P. Stone, F. G. Guinjoan, R. T. Rodrigues, H. Bruyninckx, H. Mandala, G. Christmann, J. L. Blanco-Claraco *et al.*, "Autonomous ground navigation in highly constrained spaces: Lessons learned from the second barn challenge at icra 2023 [competitions]," *IEEE Robotics & Automation Magazine*, vol. 30, no. 4, pp. 91–97, 2023.
- [24] X. Xiao, Z. Xu, A. Datar, G. Warnell, P. Stone, J. J. Damanik, J. Jung, C. A. Deresa, T. D. Huy, C. Jinyu *et al.*, "Autonomous ground navigation in highly constrained spaces: Lessons learned from the third barn challenge at icra 2024 [competitions]," *IEEE Robotics & Automation Magazine*, vol. 31, no. 3, pp. 197–204, 2024.
- [25] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," in *Conference on Robot Learning*. PMLR, 2023, pp. 2165–2183.
- [26] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi *et al.*, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, 2024.
- [27] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, " π_0 : A vision-language-action flow model for general robot control," *arXiv preprint arXiv:2410.24164*, 2024.
- [28] D. Shah, B. Osinski, S. Levine *et al.*, "Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *Conference on robot learning*. PMLR, 2023, pp. 492–504.
- [29] A.-C. Cheng, Y. Ji, Z. Yang, Z. Gongye, X. Zou, J. Kautz, E. Bryik, H. Yin, S. Liu, and X. Wang, "Navila: Legged robot vision-language-action model for navigation," *arXiv preprint arXiv:2412.04453*, 2024.
- [30] P. Ding, H. Zhao, W. Zhang, W. Song, M. Zhang, S. Huang, N. Yang, and D. Wang, "Quar-vla: Vision-language-action model for quadruped robots," in *European Conference on Computer Vision*. Springer, 2024, pp. 352–367.
- [31] N. Hirose, C. Glossop, D. Shah, and S. Levine, "OmniVla: An omni-modal vision-language-action model for robot navigation," *arXiv preprint arXiv:2509.19480*, 2025.
- [32] S. Bai, K. Chen, X. Liu, J. Wang, W. Ge, S. Song, K. Dang, P. Wang, S. Wang, J. Tang, H. Zhong, Y. Zhu, M. Yang, Z. Li, J. Wan, P. Wang, W. Ding, Z. Fu, Y. Xu, J. Ye, X. Zhang, T. Xie, Z. Cheng, H. Zhang, Z. Yang, H. Xu, and J. Lin, "Qwen2.5-vl technical report," 2025. [Online]. Available: <https://arxiv.org/abs/2502.13923>
- [33] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "Lora: Low-rank adaptation of large language models," 2021. [Online]. Available: <https://arxiv.org/abs/2106.09685>
- [34] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021. [Online]. Available: <https://arxiv.org/abs/2010.11929>
- [35] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *International conference on machine learning*. PMLR, 2018, pp. 1587–1596.
- [36] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE robotics & automation magazine*, vol. 4, no. 1, pp. 23–33, 2002.
- [37] C. Rösmann, W. Feiten, T. Wösch, F. Hoffmann, and T. Bertram, "Trajectory modification considering dynamic constraints of autonomous robots," in *ROBOTIK 2012; 7th German Conference on Robotics*. VDE, 2012, pp. 1–6.
- [38] G. Williams, A. Aldrich, and E. A. Theodorou, "Model predictive path integral control: From theory to parallel computation," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [39] Y. Lu, T. Xu, L. Wang, N. Hawes, and X. Xiao, "Decremental dynamics planning for robot navigation," *arXiv preprint arXiv:2503.20521*, 2025.

[40] OpenAI, "GPT-4o system card," 2024. [Online]. Available: <https://arxiv.org/abs/2410.21276>

[41] G. M. Kurtzer, V. Sochat, and M. W. Bauer, "Singularity: Scientific

containers for mobility of compute," *PloS one*, vol. 12, no. 5, p. e0177459, 2017.