

# RoboMatch: A Unified Mobile-Manipulation Teleoperation Platform with Auto-Matching Network Architecture for Long-Horizon Tasks

Hanyu Liu<sup>\*1</sup>, Yunsheng Ma<sup>\*1</sup>, Jiaxin Huang<sup>1</sup>, Keqiang Ren<sup>1</sup>, Jiayi Wen<sup>1</sup>, Yilin Zheng<sup>1</sup>,  
Haoru Luan<sup>1</sup>, Baishu Wan<sup>1</sup>, Pan Li<sup>1</sup>, Jiejun Hou<sup>1</sup>, Zhihua Wang<sup>1</sup>, Zhigong Song<sup>1</sup>

**Abstract**—This paper presents RoboMatch, a novel unified teleoperation platform for mobile manipulation with an auto-matching network architecture, designed to tackle long-horizon tasks in dynamic environments. Our system enhances teleoperation performance, data collection efficiency, task accuracy, and operational stability. The core of RoboMatch is a cockpit-style control interface that enables synchronous operation of the mobile base and dual arms, significantly improving control precision and data collection. Moreover, we introduce the Proprioceptive-Visual Enhanced Diffusion Policy (PVE-DP), which leverages Discrete Wavelet Transform (DWT) for multi-scale visual feature extraction and integrates high-precision IMUs at the end-effector to enrich proprioceptive feedback, substantially boosting fine manipulation performance. Furthermore, we propose an Auto-Matching Network (AMN) architecture that decomposes long-horizon tasks into logical sequences and dynamically assigns lightweight pre-trained models for distributed inference. Experimental results demonstrate that our approach improves data collection efficiency by over 20%, increases task success rates by 20–30% with PVE-DP, and enhances long-horizon inference performance by approximately 40% with AMN, offering a robust solution for complex manipulation tasks. Project website: <https://robomatch.github.io>

## I. INTRODUCTION

Imitation learning from human demonstrations offers an efficient way for robots to acquire skills, allowing humans to teach various tasks. Mobile teleoperation platforms are ideal for complex whole-body coordination tasks like opening doors or moving objects. Current systems mainly follow two approaches: decoupled platforms [1] require distinct operations for mobility and manipulation. While simplifying individual control, this divided approach may cause efficiency loss and synchronization issues. Integrated platforms face either usability limitations (e.g., Mobile ALOHA’s push-pull control [2]) or flexibility constraints (e.g., HOMIE’s fixed cockpit [3]). These limitations hinder high-quality demonstration collection and restrict imitation learning’s effectiveness in complex tasks.

Imitation learning has gained significant attention in robotics due to its practical applicability. Strategies like ACT [4] and Diffusion Policy [5] enable autonomous decision-making through end-to-end learning from demonstrations and environmental perception. Such perception includes visual [6] and proprioceptive sensing [7]. While visual

representations encode scene information, learning robust visual features for robotics remains challenging. Attention mechanisms like the Efficient Multi-scale Attention Module (EMA) [8] enhance visual perception but focus only on spatial features, ignoring frequency-domain information. Proprioceptive accuracy critically affects motion control precision. During delicate tasks like assembly or grasping, insufficient data and error accumulation often cause end-effector pose deviations, leading to failure.

As robotic intelligence advances, task planning, as a critical bridge connecting high-level instructions with low-level execution, faces growing challenges. Traditional rule-based methods often fall short in long-horizon, multi-step tasks within open environments. Emerging large language models (LLMs) [9] offer a promising alternative by parsing abstract instructions into executable task sequences through natural language understanding and logical reasoning. When integrated with visual-language models (VLMs) [10], [11], robots can transcend simple perception to achieve complex task reasoning. For instance, ReKep [12] leverages VLMs to generate physically grounded cost functions over structured 3D representations, enabling precise, constrained strategy execution. Breakthrough end-to-end vision-language-action (VLA) [13]–[15] architectures further unify perception, decision-making, and execution in a single model. However, critical challenges remain: constraint satisfaction struggles with complex multi-step tasks; VLAs exhibit limited long-horizon reasoning and hallucinations [16]; while substantial network architectures cause latency and error accumulation, hindering real-world deployment [17].

Therefore, we propose **RoboMatch**, a unified mobile-manipulation teleoperation platform with a mecha aesthetic and auto-matching network architecture for long-horizon tasks. The main contributions presented in our paper are as follows:

1) **Unified Mobile-Manipulation Teleoperation Platform:** We propose RoboMatch, a mecha-inspired platform integrating unified mobility-manipulation control, exoskeleton operation, VR-assisted perception, *etc.*, for high-quality demonstration collection.

2) **Proprioceptive-Visual Enhanced Diffusion Policy:** We design the PVE-DP framework, introducing enhancements in two key areas:

- We integrate the FE-EMA visual module (Frequency Enhanced EMA) into the Diffusion Policy visual backbone network to construct spatio-frequency-fused visual representations.

<sup>\*</sup>These authors contributed equally to this work.

Corresponding author: Zhigong Song (song\_jnu@jiangnan.edu.cn).

<sup>1</sup>Jiangsu Key Laboratory of Advanced Food Manufacturing Equipment and Technology, School of Mechanical Engineering, Jiangnan University, Wuxi 214122, China. This work was supported by the Wuxi Taihu Lake Talent Plan.

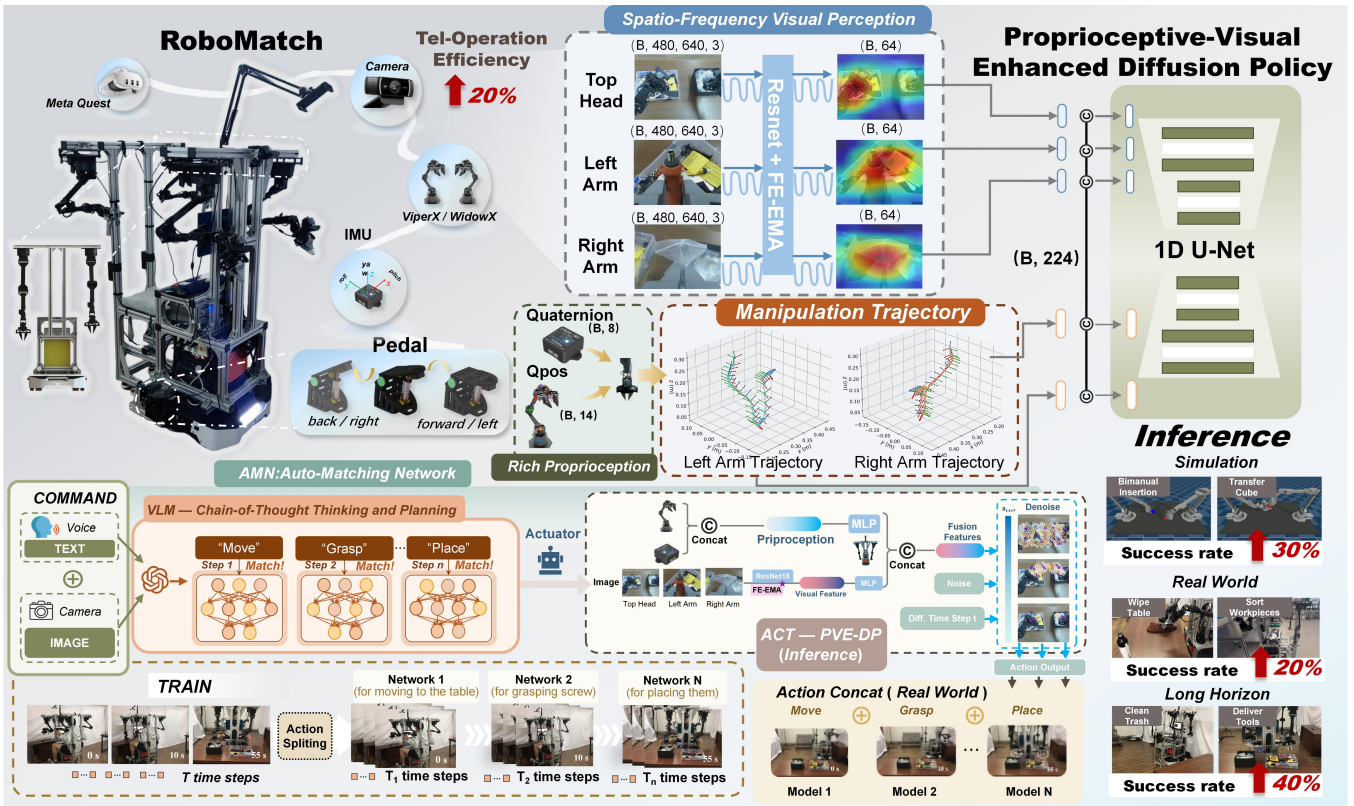


Fig. 1: Overview of the RoboMatch framework. This figure integrates three core components: (1) **RoboMatch**, a unified mobile-manipulation teleoperation platform that matches the robot base with the operation platform to achieve high-precision master-slave control and immersive observation; (2) **PVE-DP**, a policy combining spatio-frequency visual enhancement and rich proprioception to improve fine manipulation accuracy; (3) **AMN**, an architecture that combines the semantic parsing capabilities of VLMs with the execution advantages of specialized small policy networks for long-horizon execution, enabling chain-of-thought reasoning for complex task decomposition and adaptive operation.

- We mount IMU sensors on both end-effectors to accurately perceive rotational states by capturing real-time quaternion data from each arm. This information is fused with joint angles to form a comprehensive proprioceptive representation.

3) **Auto-Matching Network (AMN)**: This work introduces an AMN framework that combines VLM-based global planning with efficient compact policy execution. Through chain-of-thought reasoning, it decomposes complex tasks into hierarchical subtasks and automatically assigns them to pre-trained lightweight policy networks.

Experimental results show RoboMatch achieves 20% higher data collection efficiency. PVE-DP improves task success rates by 30% in simulation and 20% in real-world ALOHA tasks, AMN improves inference performance by around 40% in long-horizon tasks. These advancements provide an effective solution for robotic data collection, precise manipulation, and long-horizon task execution.

## II. RELATED WORK

### A. Teleoperation for Mobile Manipulation

Robot Teleoperation [18], [19] enables real-time remote control via intermediary systems and is widely used in

industrial, medical, and collaborative scenarios. However, mobile manipulation research remains nascent, especially in achieving efficient mobile base-arm coordination. While early systems adopted master-slave architectures [20], recent advances like Mobile ALOHA [2] and HOMIE [3] integrate mobility and manipulation. Some approaches leverage high-end hardware (e.g., motion capture [21], exoskeletons [3]) or adapt static interfaces (e.g., VR [22], keyboard [23]) to mobile platforms. Yet, as shown in Table I, most solutions still lack fine multi-DoF control, task generalization, operational integration, and long-term deployment capabilities, limiting their applicability in complex mobile manipulation tasks requiring high precision and integration.

### B. Imitation Learning

Imitation learning has become a key approach in robotics [28], [29], enabling robots to acquire manipulation strategies for complex tasks from expert demonstrations [30], [31]. Among these methods, behavior cloning (BC) serves as a foundational technique that directly maps observations to actions [32], offering strong transferability and proving particularly suitable for policy learning based on teleoperation data. Recent advances integrate imitation learning with

TABLE I: Comparison of Existing Robot Teleoperation Systems

Teleop System	Modality	Action Space	Cockpit	Bimanual	Exo-Teleop	Whole-Body Teleop	Wild
Mobile ALOHA [2]	Puppeteer	Joint Pos. / Base Vel.	×	✓	×	✓	✓
AirExo [24]	Puppeteer	Joint Pos.	×	✓	✓	×	✓
MoMaRT [1]	Phone	EE Pose / Base Vel.	×	×	×	✓	×
Dexcap [25]	Vision Retarget	EE Pose	×	✓	✓	×	×
MOMA-Force [26]	Kineshetic	EE Pose and Wrench	×	✓	×	×	×
UMI [27]	Kineshetic	EE Pose	×	×	×	×	✓
HOMIE [3]	Puppeteer / Pedal	Joint Pos. / Gait	Fixed	✓	✓	✓	×
<b>RoboMatch (ours)</b>	VR / Puppeteer / Pedal / Multi-View	Joint Pos. / Base Vel.	Followed	✓	✓	✓	✓

deep learning frameworks, adopting network strategies such as Transformer [4], [33] and Diffusion [5] for autonomous decision-making in complex continuous tasks. Notably, diffusion models have demonstrated remarkable generalization in handling intricate robotic tasks and are widely adopted in VLA large models [15], [34], [35].

Despite their strong task adaptability, diffusion models still exhibit noticeable limitations in motion generation precision. Cumulative errors during inference, coupled with insufficient visual perception and proprioception, significantly impede precise decision-making and efficient robot control [36].

### C. Long-Horizon Inference in Robotics

The decomposition and execution of long-horizon tasks have long been a core focus in robotics [37], with recent studies demonstrating that LLM-based planning agents offer substantial advantages in this domain [9]. The rapid development of large language models has made natural language instruction prompting a major trend in robot task planning. Recently, VLMs have demonstrated broad potential, with research mainly dividing into two technical pathways. In the first pathway, some approaches employ VLMs as task planners or formulate them as constraint satisfaction problems [10]. For instance, the ReKep system [12] uses VLMs to generate numerical cost functions defined over 3D environmental keypoints, optimizing relational constraints across the robot-object system to achieve multi-stage task execution. The second pathway explores integrated VLA models that combine VLMs with action heads, providing a feasible framework for end-to-end robot learning, as exemplified by OpenVLA [13] and  $\pi 0.5$  [14].

However, these methods still face notable limitations: constraint-based planners perform well in simple pick-and-place tasks but degrade significantly in multi-step scenarios, while VLA models suffer from inadequate long-horizon inference, frequent hallucinations, and slow inference speeds due to their large architectures [16], [17]. These issues amplify error accumulation and result in substantially lower task success rates.

## III. METHOD

As shown in Fig. 1, RoboMatch is a unified mobile-manipulation framework for data collection and policy learning, capable of directly capturing human demonstrations in real-world environments and converting them into deployable

robot policies. The system is designed with the following objectives:

**Collaborative:** RoboMatch integrates the robot and cockpit for seamless human-controlled mobile manipulation, enabling a single operator to perform grasping and mobility tasks effortlessly.

**Precise:** Our system captures end-effector (EE) rotational states and frequency-domain information from visual data, enhancing manipulation accuracy.

**Long-Horizon:** Our framework decomposes long horizon tasks into subtask actions through a chain-of-thought reasoning paradigm, dynamically assembling specialized neural networks to accomplish complex long-horizon robotic operations.

The following sections elaborate on how we achieve the above objectives through hardware design and algorithmic strategy design.

### A. RoboMatch Hardware Design and Manipulation

Our designed humanoid robot is named **RoboMatch**. The system consists of two main components: the robot body and the control cockpit. The robot body comprises two 7-DOF slave manipulators (ViperX-300), an indoor differential-drive mobile base (Tracer-2.0), three Logitech RGB cameras (C922x webcams), and two end-effector IMUs. The cockpit is equipped with two 7-DOF master arms (WidowX-250), foot pedals for base motion control, a display, keyboard, mouse, and an industrial computer. The human operator wears a VR headset (Meta Quest-3) for real-time task monitoring through multi-view visual feedback.

In terms of design, we adopted an inverted master-slave manipulator layout to form an exoskeleton-based teleoperation configuration, enhancing operational intuitiveness by aligning with the operator’s natural motion mapping. Dedicated foot pedals with embedded IMUs were developed for feet-controlled base movement. Each pedal measures rotation around the Pitch axis, mapping the angle to velocity commands: the left pedal controls forward/backward motion with speed proportional to tilt, while the right pedal enables precise left/right rotation using the same mapping logic.

As shown in Fig. 2, RoboMatch enables the operator to effortlessly perform both mobility and fine manipulation through the joint mapping of the master-slave manipulators and multi-view visual feedback. The system enables coordinated upper-lower limb movement in a unified human-robot



Fig. 2: RoboMatch teleoperation demonstration.

drive, significantly improving ergonomic adaptation and immersive control. The highly integrated hardware-software architecture not only enhances data collection efficiency but also provides high-fidelity and highly consistent human demonstration data for imitation learning and teleoperation tasks.

### B. The Model Framework of PVE-DP

As shown in Fig. 3, to enable the robot to capture richer perceptual data, we propose the **Proprioceptive-Visual Enhanced Diffusion Policy (PVE-DP)**. Building upon the original Diffusion Policy, this approach incorporates a proprioceptive enhancement module and a visual enhancement module. First, we integrate IMUs at the robot arm’s end-effectors to acquire quaternion data  $X_{\text{quat}}(w, x, y, z)$ , which satisfies the unit quaternion constraint in Eq. (1) and Eq. (2):

$$\mathbf{q} = w + x\vec{i} + y\vec{j} + z\vec{k}, \quad (1)$$

$$|\mathbf{q}|^2 = w^2 + x^2 + y^2 + z^2 = 1. \quad (2)$$

In Eq. (3), the end-effector quaternion is incorporated as a new modality and further fused with joint angle data to construct a rich proprioceptive feature representation  $P_{\text{fused}}$ :

$$P_{\text{fused}} = \text{Concat}(X_{\text{qpos}}, X_{\text{quat}}). \quad (3)$$

Although forward kinematics can compute end-effector orientation, errors in joint positions inevitably propagate to the estimated quaternion. The IMU is not intended to merely replace FK computation; instead,  $X_{\text{quat}}$  provides direct physical orientation feedback, complementing  $X_{\text{qpos}}$  at the proprioceptive level.

At the algorithmic level, we introduce a FE-EMA module and deeply integrate it with the visual backbone of the diffusion policy. This module enables efficient fusion of spatial and frequency domain features, thereby constructing spatio-frequency visual feature representations.

Given an input feature map  $X$ , its wavelet transform follows Eq. (4) via a set of scaled and translated wavelet basis functions  $\psi_{j,k}(x)$  and scaling functions  $\varphi_{j,k}(x)$ :

$$X = \sum_{j \in \mathbb{Z}} \sum_{k \in \mathbb{Z}} C_A^{j,k} \cdot \varphi_{j,k}(x) + \sum_{i=1}^3 C_D^{i,j,k} \cdot \psi_{i,j,k}(x), \quad (4)$$

where  $j$  is the scale parameter controlling frequency resolution, and  $k$  is the translation parameter governing spatial localization.  $C_A^{j,k}$  denotes the low-frequency approximation coefficients, while  $C_D^{i,j,k}$  represents the high-frequency detail coefficients capturing horizontal, vertical, and diagonal directional details.  $\mathbb{Z}$  denotes the set of integers, indicating that  $j$  and  $k$  can take any integer values.

For a 2D feature map  $X$  the wavelet transform can be implemented using 2D filters. It decomposes the feature map along rows and columns, yielding the following coefficients:  $\text{DWT}(X) = (cA, cH, cV, cD)$ , where  $cA$  denotes the low-frequency approximation coefficients, capturing global structural information;  $cH$  represents the horizontal detail components, extracting horizontal edge features;  $cV$  corresponds to the vertical detail components, capturing vertical edge features; and  $cD$  indicates the diagonal detail components, characterizing diagonal edge features.

The high-frequency feature fusion is expressed as  $F_H = cH + cV + cD$ , combined with the low-frequency feature  $cA$  to form an enhanced feature representation  $X_{\text{freq}} = \text{ResBlock}(cA + F_H)$ . Finally, the FE-EMA module generates a cross-space attention map via matrix dot-product to capture pixel-wise relationships and enhance feature representation. The spatial fusion feature  $X_{\text{space}}$  is generated through an attention mechanism and integrated with the frequency-domain feature  $X_{\text{freq}}$ , as shown in Eq. (5):

$$V_{\text{fused}} = \alpha \cdot X_{\text{space}} + (1 - \alpha) \cdot X_{\text{freq}}. \quad (5)$$

In Eq. (5),  $X_{\text{space}}$  represents the spatial domain feature generated through an attention mechanism and group normalization process. It primarily captures local regional information extracted from the input features. This feature is combined with the frequency-domain feature  $X_{\text{freq}}$  to ensure that the model captures more comprehensive multi-scale information. Such an approach enables dynamic balancing between spatial and frequency domain features.

Based on the above multi-modal fusion representation, the loss function of the diffusion policy is defined in Eq. (6):

$$\mathcal{L}(\theta) = \mathbb{E}_{t, A_t^k, \epsilon} [\|\epsilon - \epsilon_{\theta}(P_{\text{fused}}, V_{\text{fused}}, A_t^k, t)\|^2], \quad (6)$$

where  $\theta$  denotes the model parameters, and  $\epsilon_{\theta}(P_{\text{fused}}, V_{\text{fused}}, A_t^k, t)$  represents the noise predicted

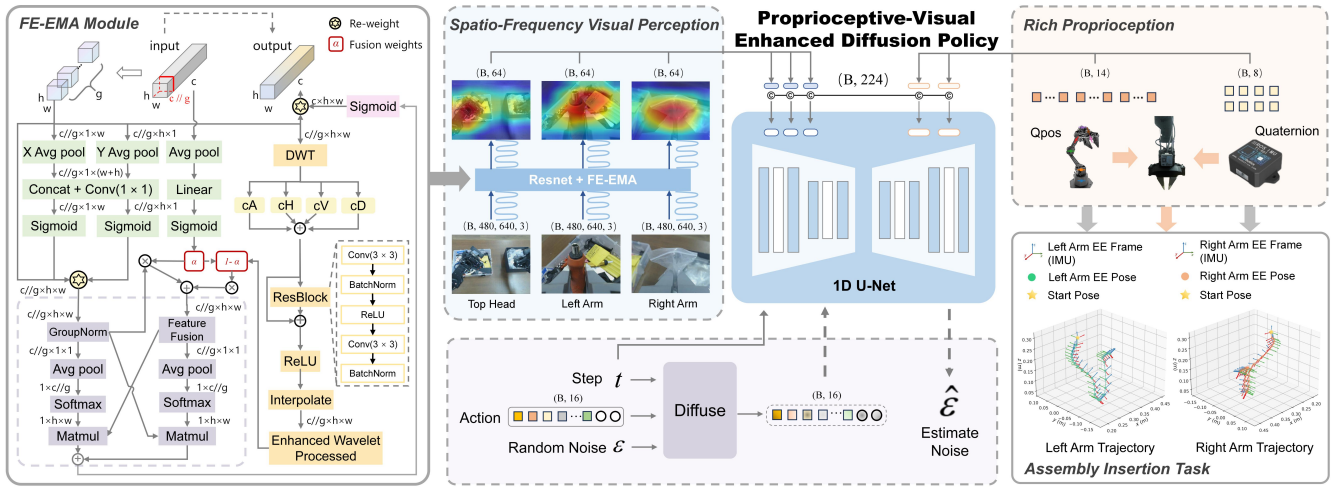


Fig. 3: Overall network architecture of PVE-DP, which consists of three core modules: (1) **FE-EMA Module** enhances robotic visual perception by integrating spatio-frequency features from visual data; (2) **Rich Proprioception** strengthens the robot’s self-state awareness through concatenating arm joint positions (Qpos) and end-effector quaternion (Quat) from IMU sensing; (3) **DP U-Net** takes the enhanced visual-proprioceptive observations as conditions and predicts fine-grained action noise to achieve precise manipulation during robot inference.

by the model based on the multi-modal fused features  $V_{\text{fused}}$  and  $P_{\text{fused}}$ , noisy action sequence  $A_t^k$ , and timestep  $t$ .

During denoising, PVE-DP continuously refines the predicted action sequence by incorporating additional wrist perception data from robotic arms IMUs, along with spatio-frequency features extracted by the FE-EMA network, enabling precise motion execution in fine manipulation tasks.

### C. Auto-Matching Network (AMN) Architecture

A critical challenge in robot imitation learning lies in long-horizon complex tasks. As execution extends, inference errors accumulate, while the absence of logical dependencies between action segments hinders a single policy from generalizing to multi-skill, multi-scenario applications.

As shown in Fig. 4, we propose the **Auto-Matching Network (AMN)** framework. This architecture integrates the vision-language model GLM-4.1V [11] and employs a “Chain of Thought” reasoning paradigm to decompose complex long-horizon tasks into logically connected subtask sequences. It dynamically matches each subtask with a pre-trained policy network for distributed inference. The core modeling concept can be expressed as Eq. (7):

$$\mathcal{T} = \{T_i\}_{i=1}^n, \quad T_i \sim \pi_i(s_t, a_t | \theta_i). \quad (7)$$

The architecture leverages the Chain of Thought cognitive paradigm to decompose a complex long-horizon task  $\mathcal{T}$  into a sequence of logically connected subtasks  $\{T_i\}$ . During training, for tasks spanning over 3000+ time steps, we introduce a step-wise saving mechanism: upon completion of each subtask  $T_i$ , the current state  $s_t$  and action  $a_t$  are saved as a subtask sample. Through this temporal decoupling strategy, the original task is decomposed into  $n$  short-horizon subtasks  $\{T_1, \dots, T_n\}$ , the  $i$ -th of which is trained with an independent network architecture  $\pi_i(\cdot | \theta_i)$ . Due to the

significant differences in perception-action patterns across subtasks, the learned network weights for each subtask exhibit high specificity.

During inference, the operator provides a command via voice or text. The robot observes the scene and sends the image and instruction to the VLM. The AMN architecture decomposes the task into sequential steps grounded in the pre-trained policy library, and then assigns each step a pre-trained specialized policy network for execution. After each subtask, the system automatically transitions to the next network until the entire task is completed.

## IV. EXPERIMENT AND ANALYSIS

This study conducts experimental validation based on the MuJoCo simulation platform and the RoboMatch real-world robot platform from Table II. On both platforms, we designed and carried out the following three verification tasks: (1) Stability of the AMN architecture in long-horizon reasoning tasks; (2) Fine manipulation capability of the PVE-DP policy in simulated and real-world tasks; (3) Statistical analysis of data collection completion rate and efficiency to evaluate RoboMatch’s applicability.

TABLE II: We train PVE-DP with short demonstrations from both simulated and real environments per task, while AMN is trained with long real-world demonstrations for each task.

Scene	No.	Task	Mode	Make-Up	Length
Sim	1	Transfer Cube	Short	100 Demos	0.2 Hrs
	2	Insertion Scripted	Short	100 Demos	0.2 Hrs
Real	3	Wipe Table	Short	50 Demos	0.3 Hrs
	4	Sort Workpiece	Short	50 Demos	0.4 Hrs
	5	Clean Trash	Long	50 Demos	1.1 Hrs
	6	Deliver Tool	Long	50 Demos	1.4 Hrs

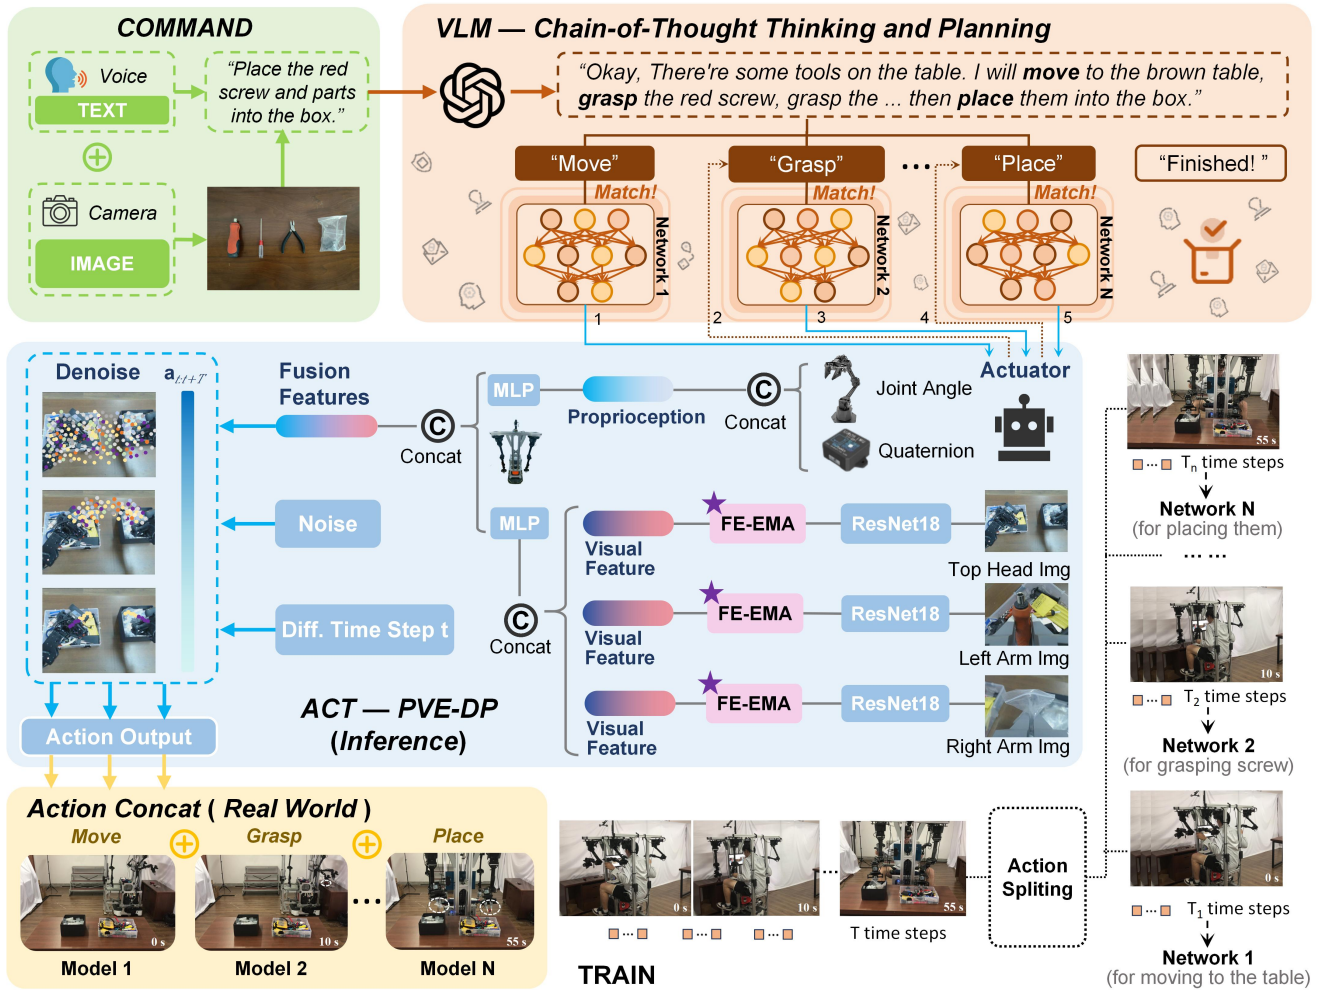


Fig. 4: Overview of AMN framework. This figure presents the Auto-Matching Network architecture, featuring: (1) **Task Decomposition** of complex task  $\mathcal{T}$  into sub-tasks for specialized policy networks; (2) **Vision-Language Input** processing both linguistic commands and visual scenes; (3) **Chain-of-Thought Thinking and Planning** that breaks tasks into sequential steps matched with pre-trained policies; (4) **Inference Execution** using PVE-DP with wrist quaternion and visual spatio-frequency fusion for precise manipulation.

### A. Experiments of AMN Architecture

As mentioned above, the AMN framework generates task plans based on the scene and instructions, decomposes them into sub-tasks, and matches each to specialized networks for long-horizon task completion. To evaluate AMN’s performance in long-horizon tasks, we compare it against baseline models including ACT, DP, and PVE-DP. Two long-horizon tasks are designed: a four-step task, *Clean Trash* (3000 time steps), and a six-step task, *Deliver Tool* (4000 time steps).

In terms of training strategy, AMN employs a phased training approach where individual subtask modules are trained independently and temporally composed during inference. In contrast, baseline methods adopt an end-to-end training strategy that directly feeds complete sequential data into the network.

As shown in Table III, AMN maintains a more stable success rate during long-horizon inference, outperforming other policies by over 40% in final task completion. While

ACT exceeds DP and PVE-DP in long-term inference, its full-task success remains below 10%, further highlighting AMN’s superior capability in complex scenarios.

### B. Ablation Experiments of PVE-DP

To thoroughly evaluate the performance of PVE-DP, we compare it against state-of-the-art baseline models in the field. The experiments include four tasks: two simulated tasks (cube transfer and bimanual insertion), and two real-robot tasks (wiping table and sorting workpiece). To facilitate algorithm validation, IMU modules were integrated into both the simulation platform and the physical RoboMatch test platform.

As shown in Table IV and Table V, our method achieves an average success rate of 75.3% without significantly increasing the number of parameters or FLOPs, outperforming all baseline methods across all four tasks. Combined with Fig. 5, these results demonstrate the superiority of the proposed




Fig. 5: Comparison between DP and the proposed PVE-DP method in simulation and real-world tasks. The first row demonstrates the inference results of DP, while the second row presents the improved performance of PVE-DP.

TABLE III: Long-horizon inference performance comparison across different methods and tasks (50 trials total). Underlined values indicate the second-highest success rates for the current task, while bolded values represent the highest success rates.

Method	Clean Trash (3000 time steps)				Deliver Tool (4000 time steps)					
	Grasp	Move	Throw	Back	Move	Grasp (screw)	Grasp (pouch)	Move	Put	Back
ACT [4]	54%	<u>46%</u>	<u>26%</u>	<u>10%</u>	100%	52%	<u>34%</u>	<u>24%</u>	<u>6%</u>	<u>2%</u>
DP U-Net [5]	30%	14%	0%	0%	100%	26%	8%	0%	0%	0%
PVE-DP (ours)	<u>66%</u>	42%	20%	8%	100%	<u>54%</u>	28%	8%	0%	0%
<b>AMN (Ours)</b>	<b>74%</b>	<b>70%</b>	<b>54%</b>	<b>50%</b>	100%	<b>68%</b>	<b>66%</b>	<b>62%</b>	<b>48%</b>	<b>44%</b>

TABLE IV: Comparison of Parameters and FLOPs.

Method	Backbone	Simulation Tasks	
		#.Param. (M)	FLOPs (B)
Diffusion Policy (DP)	ResNet18	101.7214	33.9266
DP + FE-EMA		101.7269	33.9662
DP + FE-EMA + Quat		<b>101.8416</b>	<b>33.9664</b>

TABLE V: Comparison of Success Rate (%) for Simulation and Real-World Tasks.

Method	Avg.	Task1	Task2	Task3	Task4
ACT [4]	52.0	86	32	48	42
DP U-Net [5]	43.8	80	36	39	20
NL-ACT [38]	53.8	86	40	43	46
InterACT [39]	58.8	88	46	51	50
DiT-Block [40]	56.8	89	56	47	35
DP w. EMA [8]	54.0	90	42	43	41
ACT w. quat. (ours)	63.0	96	56	54	46
DP w. quat. (ours)	61.0	98	52	51	43
DP w. FE-EMA (ours)	59.0	94	48	47	47
<b>PVE-DP (ours)</b>	<b>75.3</b>	<b>100</b>	<b>79</b>	<b>63</b>	<b>59</b>

approach in manipulation tasks. Furthermore, our method consistently exceeds baseline performance with both scripted and human demonstrations, indicating its strong capability to capture the multimodal characteristics of human data.

### C. RoboMatch Performance Experiment

To evaluate RoboMatch systematically, this study uses a split-type (locomotion and manipulation) data collection platform for baseline comparison through a dual-metric (efficiency and quality) analysis. The test set includes 100

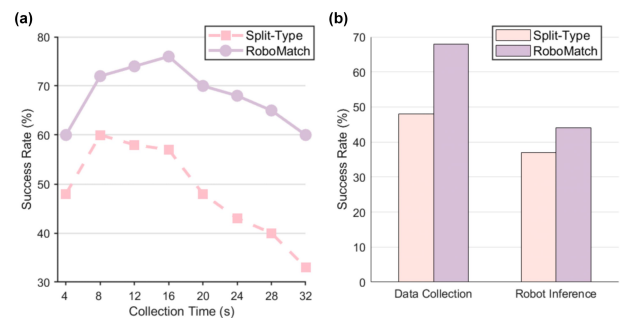


Fig. 6: The impact of RoboMatch on performance. (a) Success rate vs. task duration: Split-Type vs. RoboMatch during data collection. (b) Average success rate: Split-Type vs. RoboMatch across both data collection and inference.

demonstrations—containing real tasks from Table II and custom tasks—with duration varying from 4s to 32s at 4s intervals. We measure the one-time collection success rate for both platforms and train a unified policy on the obtained data to evaluate quality differences via inference testing.

As shown in Fig. 6, RoboMatch exhibits significantly higher collection stability with increasing task duration, consistently outperforming the decoupled platform. Quantitative results show a 20% average improvement in success rate and a 7% higher inference success rate for policies trained with RoboMatch data, demonstrating its effectiveness in enhancing robotic autonomy.

## V. CONCLUSIONS

In this paper, we propose RoboMatch, a collaborative, precise, portable, and flexible platform designed for unified whole-body mobile-manipulation teleoperation, which enables the accurate execution of operational tasks and

supports long-horizon inference in real-world environments. The proposed platform significantly reduces the complexity and human effort of teleoperation. Additionally, the PVE-DP method enhances fine manipulation capabilities, while the AMN framework improves operational flexibility, enabling the seamless integration of diverse skills and increasing success rates in long-horizon tasks. Future work will focus on optimizing the RoboMatch cockpit to enhance operator ergonomics during data collection, refining the PVE-DP method to further elevate robotic manipulation performance, and smoothing the transitions between subtasks within the AMN framework to facilitate more cohesive inference.

## REFERENCES

- [1] J. Wong *et al.*, “Error-aware imitation learning from teleoperation data for mobile manipulation,” in *Conference on Robot Learning*. PMLR, 2022, pp. 1367–1378.
- [2] Z. Fu, T. Z. Zhao, and C. Finn, “Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation,” in *Conference on Robot Learning (CoRL)*, 2024.
- [3] Q. Ben, F. Jia, J. Zeng, J. Dong, D. Lin, and J. Pang, “Homie: Humanoid loco-manipulation with isomorphic exoskeleton cockpit,” *arXiv preprint arXiv:2502.13013*, 2025.
- [4] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” *arXiv preprint arXiv:2304.13705*, 2023.
- [5] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, vol. 44, no. 10-11, pp. 1684–1704, 2025.
- [6] S. Y. Chen, “Kalman filter for robot vision: A survey,” *IEEE Transactions on Industrial Electronics*, vol. 59, no. 11, pp. 4409–4420, 2012.
- [7] L. Cappello *et al.*, “Robot-aided assessment of wrist proprioception,” *Frontiers in human neuroscience*, vol. 9, p. 198, 2015.
- [8] D. Ouyang, S. He, G. Zhang, M. Luo, H. Guo, J. Zhan, and Z. Huang, “Efficient multi-scale attention module with cross-spatial learning,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [9] B. Ichter *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Proceedings of The 6th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, K. Liu, D. Kulic, and J. Ichnowski, Eds., vol. 205. PMLR, 14–18 Dec 2023, pp. 287–318.
- [10] J. Gao, B. Sarkar, F. Xia, T. Xiao, J. Wu, B. Ichter, A. Majumdar, and D. Sadigh, “Physically grounded vision-language models for robotic manipulation,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 12 462–12 469.
- [11] W. Hong *et al.*, “Glm-4.1 v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning,” *arXiv preprint arXiv:2507.01006*, 2025.
- [12] W. Huang, C. Wang, Y. Li, R. Zhang, and L. Fei-Fei, “Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2025, pp. 4573–4602.
- [13] M. J. Kim *et al.*, “Openvla: An open-source vision-language-action model,” in *8th Annual Conference on Robot Learning*, vol. 270. PMLR, 2025, pp. 2679–2713.
- [14] P. Intelligence *et al.*, “ $\pi_0.5$ : a vision-language-action model with open-world generalization,” *arXiv preprint arXiv:2504.16054*, 2025.
- [15] J. Wen *et al.*, “Diffusionvla: Scaling robot foundation models via unified diffusion and autoregression,” in *Forty-second International Conference on Machine Learning*, 2025.
- [16] Q. Sun, P. Hong, T. D. Pala, V. Toh, U.-X. Tan, D. Ghosal, and S. Poria, “Emma-x: An embodied multimodal action model with grounded chain of thought and look-ahead spatial reasoning,” in *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2025, pp. 14 199–14 214.
- [17] R. Sapkota, Y. Cao, K. I. Roumeliotis, and M. Karkee, “Vision-language-action (vla) models: Concepts, progress, applications and challenges,” *arXiv preprint arXiv:2505.04769*, 2025.
- [18] K. Darvish *et al.*, “Teleoperation of humanoid robots: A survey,” *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1706–1727, 2023.
- [19] C. Pacchierotti and D. Prattichizzo, “Cutaneous/tactile haptic feedback in robotic teleoperation: Motivation, survey, and perspectives,” *IEEE Transactions on Robotics*, vol. 40, pp. 978–998, 2024.
- [20] J. Vertut and P. Coiffet, “Robot technology. vol. 3a. teleoperation and robotics: evolution and development,” Jan 1985.
- [21] C. Stanton, A. Bogdanovych, and E. Ratanasena, “Teleoperation of a humanoid robot using full-body motion capture, example movements, and machine learning,” in *Proc. Australasian Conference on Robotics and Automation*, vol. 8, 2012, p. 51.
- [22] A. Iyer, Z. Peng, Y. Dai, I. Guzey, S. Haldar, S. Chintala, and L. Pinto, “Open teach: A versatile teleoperation system for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2025, pp. 2372–2395.
- [23] P. Gliesche *et al.*, “Kinesthetic device vs. keyboard/mouse: a comparison in home care telemanipulation,” *Frontiers in Robotics and AI*, vol. 7, p. 561015, 2020.
- [24] H. Fang, H.-S. Fang, Y. Wang, J. Ren, J. Chen, R. Zhang, W. Wang, and C. Lu, “Airexo: Low-cost exoskeletons for learning whole-arm manipulation in the wild,” in *Proceedings of 2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 15 031–15 038.
- [25] C. Wang, H. Shi, W. Wang, R. Zhang, L. Fei-Fei, and C. K. Liu, “Dexcap: Scalable and portable mocap data collection system for dexterous manipulation,” *arXiv preprint arXiv:2403.07788*, 2024.
- [26] T. Yang, Y. Jing, H. Wu, J. Xu, K. Sima, G. Chen, Q. Sima, and T. Kong, “Moma-force: Visual-force imitation for real-world mobile manipulation,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 6847–6852.
- [27] C. Chi *et al.*, “Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots,” *arXiv preprint arXiv:2402.10329*, 2024.
- [28] S. Schaal, “Is imitation learning the route to humanoid robots?” *Trends in cognitive sciences*, vol. 3, no. 6, pp. 233–242, 1999.
- [29] E. Jang *et al.*, “Bc-z: Zero-shot task generalization with robotic imitation learning,” in *Conference on Robot Learning*. PMLR, 2022, pp. 991–1002.
- [30] H. Kim, Y. Ohmura, and Y. Kuniyoshi, “Goal-conditioned dual-action imitation learning for dexterous dual-arm robot manipulation,” *IEEE Transactions on Robotics*, vol. 40, pp. 2287–2305, 2024.
- [31] J. Gao, X. Jin, F. Krebs, N. Jaquier, and T. Asfour, “Bi-kvil: Keypoints-based visual imitation learning of bimanual manipulation tasks,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16 850–16 857.
- [32] F. Torabi, G. Warnell, and P. Stone, “Behavioral cloning from observation,” *arXiv preprint arXiv:1805.01954*, 2018.
- [33] H. Kim, Y. Ohmura, and Y. Kuniyoshi, “Transformer-based deep imitation learning for dual-arm robot manipulation,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8965–8972.
- [34] J. Liu *et al.*, “Hybridvla: Collaborative diffusion and autoregression in a unified vision-language-action model,” *arXiv preprint arXiv:2503.10631*, 2025.
- [35] H. Zhen, X. Qiu, P. Chen, J. Yang, X. Yan, Y. Du, Y. Hong, and C. Gan, “3d-vla: a 3d vision-language-action generative world model,” in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 61 229–61 245.
- [36] R. Wolf, Y. Shi, S. Liu, and R. Rayyes, “Diffusion models for robotic manipulation: A survey,” *Frontiers in Robotics and AI*, vol. 12, p. 1606247, 2025.
- [37] G. Konidaris, S. Kuindersma, R. Grupen, and A. Barto, “Robot learning from demonstration by constructing skill trees,” *The International Journal of Robotics Research*, vol. 31, no. 3, pp. 360–375, 2012.
- [38] K. Rohling, “Integrating natural language instructions into the action chunking transformer for multi-task robotic manipulation,” <https://github.com/krohling/nl-act>, 2024.
- [39] A. C.-W. Lee, I. Chuang, L.-Y. Chen, and I. Soltani, “Interact: Inter-dependency aware action chunking with hierarchical attention transformers for bimanual manipulation,” in *Conference on Robot Learning*, vol. 270. PMLR, 2025, pp. 1730–1743.
- [40] S. Dasari, O. Mees, S. Zhao, M. K. Srirama, and S. Levine, “The ingredients for robotic diffusion transformers,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 15 617–15 625.