

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

# Neural Internal Model Control: Learning a Robust Control Policy via Predictive Error Feedback

Feng Gao<sup>1</sup>, Chao Yu<sup>13†</sup>, Yu Wang<sup>1</sup>, Yi Wu<sup>12†</sup>

**Abstract**—Accurate motion control in the face of disturbances within complex environments remains a major challenge in robotics. Classical model-based approaches often struggle with nonlinearities and unstructured disturbances, while reinforcement learning (RL)-based methods can be fragile when encountering unseen scenarios. In this paper, we propose a novel framework, Neural Internal Model Control (NeuralIMC), which integrates model-based control with RL-based control to enhance robustness. Our framework streamlines the predictive model by applying Newton-Euler equations for rigid-body dynamics, eliminating the need to capture complex high-dimensional nonlinearities. This internal model combines model-free RL algorithms with predictive error feedback. Such a design enables a closed-loop control structure to enhance the robustness and generalizability of the control system. We demonstrate the effectiveness of our framework on both quadrotors and quadrupedal robots, achieving superior performance compared to state-of-the-art methods. Furthermore, real-world deployment on a quadrotor with rope-suspended payloads highlights the framework’s robustness in sim-to-real transfer. Our code is released at <https://github.com/thu-uav/NeuralIMC>.

**Index Terms**—Robust Control, Reinforcement Learning, Sensorimotor Learning

## I. INTRODUCTION

Achieving stable and precise motion in complex, disturbed environments is essential for robots, such as quadrotors maintaining stable flight in wind fields and legged robots carrying loads over diverse terrains. This challenge, known as the robust control or adaptive control problem, has been extensively studied [1]. Classical adaptive control methods adjust control parameters in real-time to compensate for dynamic environmental changes, often using mathematical models of the system and disturbances [2], [3], [4], [5], while robust control methods utilize more informative signals to enhance robustness, like Internal Model Control [6]. Despite their guaranteed stability, these methods rely on accurately modeling system dynamics, which is challenging for nonlinear systems with unstructured disturbances. In contrast, reinforcement learning (RL) methods have shown remarkable success in continuous control across

Manuscript received: November 11, 2024; Revised March 5, 2025; Accepted May 1, 2025.

This paper was recommended for publication by Editor Jens Kober upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by National Natural Science Foundation of China (No.62406159, 62325405), Postdoctoral Fellowship Program of CPSF under Grant Number (GZC20240830, 2024M761676), China Postdoctoral Science Special Foundation 2024T170496.

<sup>†</sup> Corresponding authors. {zoeyuchao, jxwuyi}@gmail.com.

<sup>1</sup> Tsinghua University, Beijing, China.

<sup>2</sup> Shanghai Qi Zhi Institute, Shanghai, China.

<sup>3</sup> Beijing Zhongguancun Academy, Beijing, China.

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

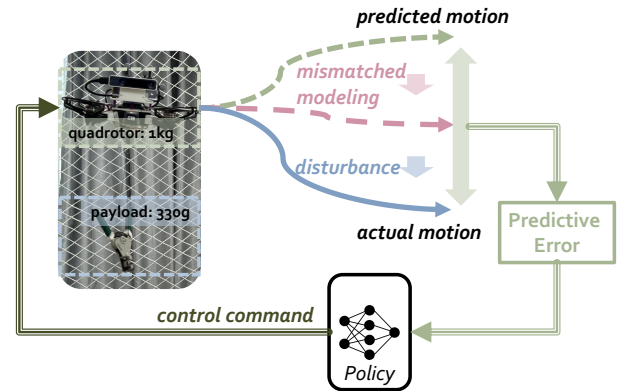


Fig. 1: Predictive errors can be caused by mismatched system modeling and/or external disturbances.

various robot morphologies, including skill learning for armed robots [7], [8], high-speed quadrotor racing [9], [10], and adaptive locomotion for legged robots [11], [12], [13], [14], [15]. Large-scale training with randomized dynamics and environmental parameters enables RL-based control methods to generalize across different conditions and environments [16], [17], allowing robots to perform reliably even in previously unseen and highly varying scenarios.

There has been a recent trend towards integrating model-based control with RL-based neural control to harness the strengths of both worlds [18], [19], [13], [20]. One representative framework is to use predicted future states or estimated disturbances as an additional input signal to obtain a more informative model-free RL policy. Long et al. [13] employ unsupervised contrastive representation learning to train a predictive encoder capable of estimating the next state in the latent space. Conversely, Huang et al. [20] use a handcrafted  $\mathcal{L}_1$  adaptive law [21] to explicitly estimate applied force disturbances, resulting in a robust trajectory tracking controller that can handle unexpected wind fields. Although promising, integrating model-based and RL-based controls faces significant challenges. Predictive models often struggle with variability when training data cannot cover real-world dynamics, and handcrafted disturbance estimators require precise tuning, limiting their feasibility for complex systems and adaptability to dynamic changes. For establishing a more general method, we need to develop a structure that effectively combines these controls to fully leverage their advantages and address these limitations.

In this paper, we propose a new learning-based control methods which leverages a predictive error feedback to en-

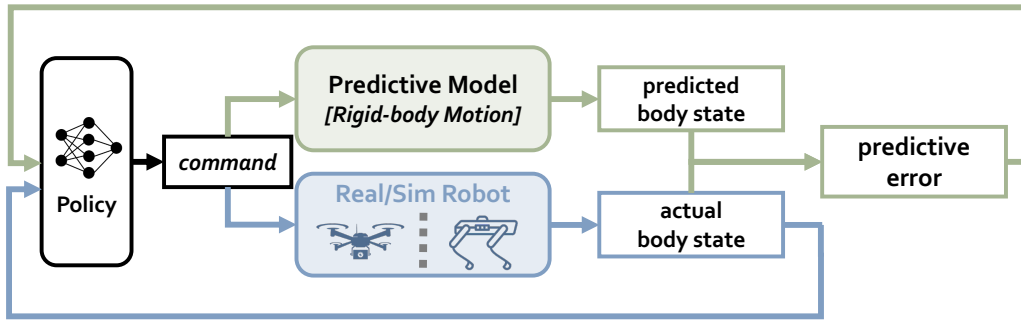


Fig. 2: The closed-loop NeuralIMC framework using an explicit predictive model to calculate the next body state, applicable to both quadrotors and legged robots.

enhance control robustness, inspired by model-based control. By simplifying the model to focus on the robot’s main body and treating its dynamics as a rigid transformation, we bypass the high-dimensional nonlinearity of complex robot systems and environmental interactions. This simplification enables us to use an explicit predictive model without requiring parameter tuning, establishing a foundation for robust closed-loop control. We compute the predictive error—defined as the discrepancy between plant and model responses—as a feedback signal to the policy, enhancing robustness and generalizability through informative feedback and extensive training. Our framework, shown in Fig. 2, can be easily integrated with other learning algorithms by incorporating this feedback loop. To demonstrate the effectiveness and versatility of our method, we apply it to two prevalent robot morphologies: quadrotors, which demand precise 3D control, and quadrupedal robots, which require adaptive locomotion to navigate complex terrain interactions. Empirical results indicate that our framework can surpass recent state-of-the-art control methods in both domains. Furthermore, we conduct a real-world deployment on a quadrotor with various rope-suspended payloads to illustrate its robustness in sim-to-real transfer.

## II. RELATED WORK

### A. Classical control and disturbance estimation

Classical adaptive controllers require estimating system parameters in a closed-loop manner to adjust control actions in real-time.  $\mathcal{L}_1$  adaptive control has been widely used to estimate force disturbances in quadrotors [22], [23]. Neural-Fly [18] incorporated pre-trained representations to enhance rapid online adaptation through deep learning, achieving stable flight in previously unseen wind conditions. Internal Model Control (IMC) [6], [3] is a general structure for robust control that uses a predefined model to predict system responses and estimate disturbances. However, applying IMC to more complex and nonlinear systems is challenging due to the difficulty of crafting reliable models for such systems. Hybrid Internal Model (HIM) further leverages a learning-based internal model that employs unsupervised contrastive learning, with the learned model acting as a predictive encoder for estimating system responses in the latent space, as proposed by Long et al. [13]. Despite their advantages, these methods either rely on explicit disturbance estimation or are heavily

influenced by prediction accuracy, which can significantly affect performance if compromised.

### B. Learning-based adaptive control

With the rapid advancement of highly parallel simulators [17], [24], [25], reinforcement learning (RL) methods are increasingly demonstrating powerful capabilities to achieve robust control, including quadrotors [26], [20] and legged robots [11], [12], [13], [14], [15]. Unlike classical model-based control, these learning-based methods often use model-free RL algorithms like PPO [27] to train neural controllers in large-scale randomized environments. Koryakovskii et al. [28] proposed to incorporate an RL-based policy to generate residual actions to compensate nonlinear model predictive control in model-mismatch cases. Rapid Motor Adaptation (RMA) [12] is a pure RL-based method, which predicts environmental parameters in the latent space using state-action histories through a teacher-student training procedure. This approach has been extended to quadrotors [26], bipedal robots [29], and humanoid robots [15]. A follow-up study found that a dual-history architecture, leveraging both short-term and long-term input/output history, achieves superior performance with a straightforward one-stage training process [14]. Conversely, DATT [20] incorporates an  $\mathcal{L}_1$  adaptation law [21] to estimate forces applied to quadrotors directly, which are then used as policy inputs. CAL [28] integrates RL with nonlinear model predictive control (NMPC) by employing RL to generate residual actions that augment NMPC’s outputs. Although both DATT and CAL have proven effective with precise parameter tuning, they remain sensitive to handcrafted settings in their classical control components, limiting their generalizability to more complex systems.

## III. METHODS

In this section, we present a general framework that bridges model-based and RL-based control through a simplified predictive model, named NeuralIMC. We start by introducing the basics of Internal Model Control (IMC) and our extension to integrate it with RL-based neural control (Sec. III-A). By leveraging the model-free RL algorithm PPO [27], we reduce the complexity and accuracy requirements typically associated with classical model-based control. Next, we detail how rigid transformation is used to predict the robot’s body

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

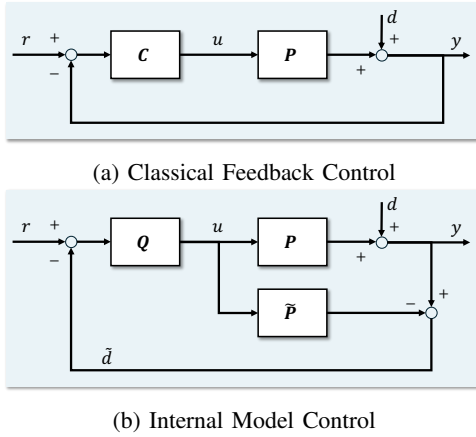


Fig. 3: Feedback structures of classical feedback control and Internal Model Control (IMC).

state, establishing a general approach to estimate disturbances and incorporate them into the control policy (Sec. III-B). Finally, we explain the application of this framework to two specific robot morphologies: quadrotors and quadrupedal robots (Sec. III-C).

#### A. Internal Model Control

Internal Model Control (IMC) is a robust control framework that outperforms classical feedback control in handling model uncertainties and disturbances. Unlike feedback control, which corrects errors in real time, IMC uses a predictive model to anticipate and compensate for disturbances. Fig. 3 illustrates their structures, where  $P$  and  $\tilde{P}$  are the plant and internal models,  $Q$  is the IMC controller, and  $C$  is the classical feedback controller. The reference, disturbance, and output are denoted by  $r$ ,  $d$ , and  $y$ , respectively. IMC computes the difference  $\tilde{d} = \tilde{y} - y$  between predicted and actual responses, adjusting control inputs accordingly for enhanced robustness.

IMC is straightforward when the plant is easily modeled [3] but becomes challenging for nonlinear systems where approximations fail to capture true dynamics [30], [31]. This reliance on precise models limits its use in robotics, such as quadrotors with suspended payloads or legged robots on uneven terrain, where accurate system modeling is difficult. Our method integrates IMC with large-scale RL training to address these challenges. A simplified predictive model provides a rough yet stable estimate, while large-scale randomized simulations train a robust control policy that compensates for modeling inaccuracies. RL optimizes the policy with the predictive error feedback, revealing external disturbances and nonlinear dynamics, leading to more adaptive and effective control in real-world settings.

#### B. Design of a Simplified Predictive Model

The predictive model is key to our framework, enabling discrepancy detection and compensation. Its design faces two challenges: accurately estimating system response and closely aligning predictions with actual behavior. Precise estimation

ensures the feedback loop reflects the true robot state, while alignment ensures reliability under varying conditions.

To streamline state estimation across diverse robots, we model only the rigid body state, excluding actuator- or limb-specific states. For robots equipped with monocular or binocular cameras and built-in IMUs, numerous visual(-inertial) odometry methods provide robust state estimations [32], [33], [34], [35], which can be used as the actual system response within the IMC structure. By treating the robot as a rigid body and using Euclidean transformation to estimate the next state, we significantly reduce the modeling complexity, making it more practical and efficient.

Here, we first present a general approach for formulating the predictive model and calculating the error, followed by specific adaptations for different robot morphologies in the next section. We assume that the model inputs include velocity or acceleration commands, and define the current body state as  $\mathbf{x}_k = [\mathbf{p}_k, \mathbf{q}_k, \mathbf{v}_k, \boldsymbol{\omega}_k] \in \mathbb{R}^{13}$ , where  $\mathbf{p}_k$  is the position,  $\mathbf{q}_k$  is the quaternion (orientation),  $\mathbf{v}_k$  is the velocity, and  $\boldsymbol{\omega}_k$  is the angular velocity. The time interval between two steps is denoted as  $\Delta t$ .

The next position and velocity of the body can be computed via Newton's law [36]:

$$\tilde{\mathbf{p}}_{k+1} = \mathbf{p}_k + \mathbf{v}_k \Delta t + \frac{1}{2} \mathbf{a}_k \Delta t^2, \quad (1)$$

$$\tilde{\mathbf{v}}_{k+1} = \mathbf{v}_k + \mathbf{a}_k \Delta t. \quad (2)$$

For rotation, we adopt a discrete update rule while handling singularities. First, we compute the change in angle  $\Delta \theta$  with the angular velocity  $\boldsymbol{\omega}_k$  and angular acceleration  $\dot{\boldsymbol{\omega}}_k$ :

$$\Delta \theta = \boldsymbol{\omega}_k \Delta t + \frac{1}{2} \dot{\boldsymbol{\omega}}_k \Delta t^2. \quad (3)$$

Next, we follow [37] to apply the exponential map to project the rotation from  $\mathbb{R}^3$  to  $SO(3)$ , obtaining the corresponding quaternion  $\mathbf{q}_{\Delta \theta}$ :

$$\mathbf{q}_{\Delta \theta} = \begin{cases} \left[ \cos\left(\frac{\|\Delta \theta\|}{2}\right), \frac{\sin\left(\frac{\|\Delta \theta\|}{2}\right)}{\|\Delta \theta\|} \cdot \Delta \theta \right], & \|\Delta \theta\| \leq \sqrt[4]{\epsilon_{\text{machine}}} \\ \left[ \cos\left(\frac{\|\Delta \theta\|}{2}\right), \left(\frac{1}{2} + \frac{\|\Delta \theta\|^2}{48}\right) \cdot \Delta \theta \right], & \text{otherwise} \end{cases} \quad (4)$$

Finally, we can predict the next quaternion and angular velocity as follows, where  $\otimes$  denotes quaternion multiplication:

$$\tilde{\mathbf{q}}_{k+1} = \mathbf{q}_k \otimes \mathbf{q}_{\Delta \theta} \quad (5)$$

$$\tilde{\boldsymbol{\omega}}_{k+1} = \boldsymbol{\omega}_k + \dot{\boldsymbol{\omega}}_k \Delta t \quad (6)$$

We have established an explicit method to predict the next body state, assuming the required motion can be represented by commands at the acceleration or velocity level. Then, given the predicted body state  $\tilde{\mathbf{x}}_{k+1} = [\tilde{\mathbf{p}}_{k+1}, \tilde{\mathbf{q}}_{k+1}, \tilde{\mathbf{v}}_{k+1}, \tilde{\boldsymbol{\omega}}_{k+1}]$  and the actual body state  $\mathbf{x}_{k+1}$ , we compute the predictive error  $\tilde{\mathbf{d}}_{k+1}$ , which is then fed into the policy to provide effective feedback signals.

$$\tilde{\mathbf{d}}_{k+1} = \left[ \Delta \tilde{\mathbf{p}}_{k+1}, 1 - \cos \tilde{\Theta}, \Delta \tilde{\mathbf{v}}_{k+1}, \Delta \tilde{\boldsymbol{\omega}}_{k+1} \right] \quad (7)$$

Here,  $\Delta \tilde{\mathbf{p}}_{k+1}$ ,  $\Delta \tilde{\mathbf{v}}_{k+1}$ , and  $\Delta \tilde{\boldsymbol{\omega}}_{k+1}$  represent the differences in position, velocity, and angular velocity, respectively. We

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

TABLE I: Parameter ranges for dynamics randomization on quadrotor.  $K_p^0 = [37, 37, 11]$

Parameter Name	Train Range	Eval Range
Mass ( $M$ )	[0.142, 0.950]	[0.114, 1.140]
Arm Length ( $L$ )	[0.046, 0.200]	[0.037, 0.240]
Width-to- $L$ Ratio ( $\alpha_w$ )	[1.0, 1.414]	[0.852, 1.732]
Height-to- $L$ Ratio ( $\alpha_h$ )	[0.577, 1.0]	[0.268, 1.414]
Motor Force Constant ( $K_f$ )	[1.15e-7, 7.64e-6]	[9.16e-8, 9.17e-6]
Max Thrust-to-Weight Ratio	[2.0, 3.5]	[1.6, 4.2]
Motor Drag Constant ( $\kappa$ )	[0.0041, 0.0168]	[0.0033, 0.0201]
Motor Time Constant ( $\tau_c$ )	[0.3, 0.5]	[0.24, 0.6]
P Gains ( $K_p$ )	$[-0.3, 0.3] \times K_p^0$	$[-0.5, 0.5] \times K_p^0$

use the cosine distance to quantify the difference between orientations, where  $\tilde{\Theta}$  is the angle difference between the two quaternions.

### C. Specifications for Common Robots

In this section, we further explain how to derive velocity or acceleration commands for various robots, including quadrotors and quadrupedal robots. We demonstrate how these commands align with existing control methods, ensuring seamless integration into the policy learning process.

1) *Quadrotors*: Traditional quadrotor controllers use cascaded structures [38], allowing flexible integration of neural controllers. Recent studies show that using collective thrust and bodyrate (CTBR) as the action space enhances resilience to dynamics mismatches, improves domain transfer, and boosts agility [39], [20], [9]. This design aligns with our framework: the policy’s bodyrate command updates rotation, while collective thrust determines linear acceleration. Consequently, predictive errors, computed from prior policy output and successive body states, integrate directly into the current policy input without modifying the learning algorithm.

2) *Quadrupedal robots*: Unlike quadrotors, legged robots, particularly quadrupeds and bipeds, are well-suited for RL-based control due to their high-dimensional actuation and terrain interactions [8], [11], [12], [14], [19]. Neural controllers typically generate actuator targets to achieve desired body velocities ( $v_x, v_y, \omega_z$ ). To integrate our predictive model into the control loop, we use reference velocity commands as model inputs, deriving linear and angular accelerations by differentiating consecutive commands. Setting other accelerations to zero ensures seamless integration with existing policies. This method is compatible with other algorithms via a plug-and-play feedback loop augmenting policy inputs.

## IV. EXPERIMENTAL RESULTS

To evaluate our framework’s performance and versatility, we conduct experiments on quadrotors and quadrupedal robots, comparing against state-of-the-art methods in each domain. Then, we perform ablation studies on the feedback structure and noise tolerance of the simplified predictive model. Finally, we present sim-to-real deployment results on a quadrotor with two rope-suspended payloads, showcasing its effectiveness in enhancing system robustness.

### A. Experiment Settings

1) *Quadrotors*: For quadrotors, controllers are trained using a customized simulator based on the dynamics shown in [40], where the collective thrust and bodyrate commands are followed by a PD controller. The fourth-order Runge-Kutta method [41] updates quadrotor states at each simulation step with a update frequency at 200 Hz, and the controller’s execution frequency is 50 Hz. Following Huang et al. [20], we train controllers on two types of reference trajectories, including smooth chained polynomial trajectories and infeasible zigzag trajectories. External disturbances are applied as randomized force perturbations within  $[-3.3 \text{ m/s}^2, 3.3 \text{ m/s}^2]$ , representing Brownian motion variations as in [20]. Dynamical parameters are randomized at the start of each episode for training, and unseen parameters are used for evaluation to test robustness. The detailed randomization range is shown in Tab. I. We train and evaluate each method using five random seeds, reporting the average performance across 1,024 parallel trials. For evaluation, we test controllers trained with smooth trajectories on two types of randomly generated smooth paths: circle and chained polynomial, while controllers trained on infeasible paths are tested on 5-point star and zigzag trajectories.

2) *Quadrupedal robots*: Following Long et al. [13], we train controllers for the Unitree Aliengo using Isaac Gym [17] with 4096 parallel environments and a 100-step rollout. The training curriculum includes progressively challenging terrains and reference commands. Controllers are trained on terrains such as pyramid slopes, rough pyramid slopes, stairs, and flats with discrete obstacles, with distributions of [0.1, 0.2, 0.6, 0.1]. We evaluate methods by tracking x-y linear velocity or yaw rate commands under two conditions: no disturbances and random force pushes. The x-y linear velocity command is sampled from  $[-1.0, 1.0]$  m/s, and the yaw rate command from  $[-1.0, 1.0]$  rad/s. During linear velocity tracking evaluations, the angular velocity target is set to zero, and vice versa for angular velocity tracking. For linear velocity tracking, the forward and lateral velocities are uniformly sampled from  $[-1 \text{ m/s}, 1 \text{ m/s}]$ . For angular velocity tracking tasks, the yaw rate is uniformly sampled from  $[-1 \text{ rad/s}, 1 \text{ rad/s}]$ . The tracking error is averaged over five random seeds, each with 4096 parallel environments. To validate robustness, we emulate an impulse push by setting a randomized base velocity to the robot every 16 steps, uniformly sampled within the maximum range. The maximum impulse velocity during training is 1 m/s, increased to 1.5 m/s for evaluation.

### B. Experiments on Quadrotors

We compare our method for quadrotors against two classical control methods ( $\mathcal{L}_1$ -PID and  $\mathcal{L}_1$ -MPPI) and four learning-based baselines (CAL [28], PPO [27], RMA [12], DATT [20]), using two variants each for RMA and DATT to enable deeper analyses. All methods use the same PPO training parameters. For CAL [28], a sum combination of MPPI and RL, MPPI serves as a sampling-based non-linear MPC method, with the RL policy compensating the MPPI output. We set the sampling number to 4096 and the rollout horizon to 20 in MPPI’s sampling procedure. PPO employs a three-layer MLP policy

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

TABLE II: Simulation results for tracking errors (unit: m) with quadrotors, averaged over five random seeds. **Bold values** denote the best results. **Blue cells** indicate performance under random external force disturbances, with percentages in brackets representing the relative change in error compared to static conditions. T and S refer to the teacher and student models for RMA, respectively, while DATT with GT indicates the use of ground-truth disturbance inputs instead of  $\mathcal{L}_1$  estimates.

Config	Method	Smooth Trajectory				Infeasible Trajectory			
		Circle		Chained Poly		5-Point Star		Zigzag	
Train	$\mathcal{L}_1$ -PID	0.214	0.237 (+10.74%)	0.043	0.054 (+25.58%)	<i>crashed</i>		<i>crashed</i>	
	$\mathcal{L}_1$ -MPPI	0.202	0.215 (+6.44%)	0.092	0.117 (+27.17%)	0.193	0.216 (+11.91%)	0.169	0.186 (+10.06%)
	CAL	0.087	0.157 (+80.46%)	0.056	0.118 (+110.71%)	0.103	0.175 (+69.90%)	0.105	0.164 (+56.19%)
	PPO	0.102	0.169 (+65.69%)	0.065	0.117 (+80.00%)	0.097	0.172 (+77.32%)	0.095	0.154 (+62.11%)
	RMA (T)	<b>0.071</b>	<b>0.080</b> (+12.68%)	0.043	0.047 (+9.30%)	0.087	0.094 ( <b>+8.05%</b> )	0.074	0.082 (+10.81%)
	RMA (S)	0.084	0.110 (+30.95%)	0.048	0.064 (+33.33%)	0.116	0.157 (+35.34%)	0.112	0.142 (+26.79%)
	DATT ( $\mathcal{L}_1$ )	0.084	0.128 (+52.38%)	0.051	0.076 (+49.02%)	0.089	0.163 (+83.15%)	0.077	0.138 (+79.22%)
	DATT (GT)	0.084	0.093 ( <b>+10.71%</b> )	0.052	0.055 ( <b>+5.77%</b> )	0.089	0.098 (+10.11%)	0.077	0.084 ( <b>+9.09%</b> )
	Ours	0.072	0.083 (+15.28%)	<b>0.032</b>	<b>0.038</b> (+18.75%)	<b>0.065</b>	<b>0.073</b> (+12.31%)	<b>0.057</b>	<b>0.064</b> (+12.28%)
	Eval	$\mathcal{L}_1$ -PID	0.372	0.416 (+11.82%)	0.061	0.107 (+75.41%)	<i>crashed</i>		<i>crashed</i>
$\mathcal{L}_1$ -MPPI		0.212	0.234 (+10.38%)	0.097	0.132 (+36.08%)	0.207	0.253 (+22.22%)	0.176	0.218 (+23.86%)
CAL		0.224	0.321 (+43.30%)	0.171	0.251 (+46.78%)	0.262	0.364 (+38.93%)	0.265	0.358 (+35.09%)
PPO		0.265	0.351 (+32.45%)	0.183	0.273 (+49.18%)	0.257	0.374 (+45.53%)	0.265	0.367 (+38.49%)
RMA (T)		0.117	0.147 (+25.64%)	0.072	0.097 (+34.72%)	0.151	0.192 (+27.15%)	0.145	0.187 (+28.97%)
RMA (S)		0.280	0.360 (+28.57%)	0.174	0.243 (+39.66%)	0.358	0.424 (+18.44%)	0.349	0.417 ( <b>+19.48%</b> )
DATT ( $\mathcal{L}_1$ )		0.137	0.523 (+281.75%)	0.094	0.354 (+276.60%)	0.169	0.585 (+246.15%)	0.160	0.559 (+249.38%)
DATT (GT)		0.139	0.167 (+20.14%)	0.094	0.119 ( <b>+26.60%</b> )	0.178	0.197 ( <b>+10.67%</b> )	0.156	0.189 (+21.15%)
Ours		<b>0.113</b>	<b>0.135</b> (+19.47%)	<b>0.054</b>	<b>0.073</b> (+35.19%)	<b>0.103</b>	<b>0.134</b> (+30.10%)	<b>0.101</b>	<b>0.132</b> (+30.69%)

and a temporal convolutional network (TCN) to process future waypoints in the reference trajectory, which is inherited by all other methods. In addition, RMA adopts a teacher-student training framework, with RMA(T) as the teacher model and RMA(S) as the student. RMA(T) integrates dynamic parameters and ground-truth force disturbances, while RMA(S) uses a TCN to approximate the teacher’s latent state from state-action histories. In contrast, DATT utilizes an  $\mathcal{L}_1$  disturbance estimator and tracking error inputs to improve trajectory-tracking robustness. To separate the  $\mathcal{L}_1$  estimator’s effect, we also include a privileged variant, DATT(GT), which directly uses ground-truth force disturbances in place of the estimator.

As shown in Tab. II, our method consistently outperforms others across nearly all tasks, particularly under disturbances and out-of-distribution (OOD) system dynamics.

Among estimator-based baselines—both explicit (e.g.,  $\mathcal{L}_1$  estimator in  $\mathcal{L}_1$ -PID,  $\mathcal{L}_1$ -MPPI, DATT) and implicit (e.g., RMA’s neural encoder)—DATT is the most resilient, benefiting from its explicit  $\mathcal{L}_1$  estimator coupled with an RL-based controller. Nevertheless, all estimator-based methods degrade sharply under our extreme out-of-distribution scenarios—where internal dynamics shift or sudden external forces occur—because their estimators become unreliable and control robustness is compromised. Estimator-free approaches such as CAL and PPO also fail to manage these varied disturbances. In contrast, our method computes the predictive error between expected and observed body transformations, directly signaling state drift and disturbance magnitude. A policy trained on this error maintains strong performance across both nominal and unexpected conditions.

We also acknowledge that limited hyperparameter sweeps in our experiments may contribute to the baselines’ degraded performance. To ensure a fair comparison, we include privileged variants (DATT (GT) and RMA (T)) that receive ground-truth states rather than estimates, representing an upper tuning

bound. Even against these privileged methods, our predictive error-driven controller achieves superior and more stable control, supporting our hypothesis that explicit predictive error feedback offers a simple, generalizable path to robust performance.

### C. Experiments on Quadrupedal Robots

For quadrupedal robots, we compare our methods with PPO [27], RMA [12] and HIM [13]. PPO and RMA serve as standard and common baselines in continuous control for quadrupeds, while HIM leverages unsupervised contrastive learning to train a predictive encoder for the next latent state and body velocity. Here, we directly use the teacher-policy of RMA to set a strong baseline. Our method enhances the PPO policy by integrating predictive error feedback. Additionally, we evaluate a variant of our method, adding predictive error feedback to both the inputs of the HIM policy and its predictive encoder during training and evaluation.

As shown in Tab. III, our method outperforms HIM by 26–44% across various test conditions and matches or surpasses the privileged RMA. The combination of our framework with HIM also shows improvements over the original HIM. Interestingly, the predictive encoder in HIM shows a slight degradation compared to standard PPO when combined with our NeuralIMC structure. We hypothesize that this is because the predictive error provides sufficient signal for robust control, rendering the additional impact of predictive encoding negligible. This suggests that predictive error feedback alone is sufficiently informative for enhancing control robustness, simplifying the neural control pipeline, and facilitating integration with other learning algorithms.


### D. Ablation Studies

In this section, we conduct three ablation studies to analyze our framework. Inspired by [14], we first examine whether


**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

TABLE III: Simulation results on quadrupedal robots, averaged over five random seeds. The tracking error for linear velocity and angular velocity are computed as  $\|v_{x,y} - v_{x,y}^{\text{target}}\|_2$  (unit: m/s) and  $\|\omega_{\text{yaw}} - \omega_{\text{yaw}}^{\text{target}}\|_2$  (unit: rad/s), respectively. The best results for each case are highlighted in **bold**.

Disturbance	Method	Slopes		Rough Slopes		Stairs		Discrete Obstacles	
		Linear	Angular	Linear	Angular	Linear	Angular	Linear	Angular
None	PPO	0.132	0.085	0.143	0.087	0.200	0.085	0.163	0.085
	RMA	0.084	0.050	0.086	0.051	0.127	0.050	0.110	0.050
	HIM	0.100	0.067	0.103	0.068	0.181	0.067	0.147	0.067
	Ours (w/ HIM)	0.077	0.052	0.079	0.053	0.126	0.053	0.105	0.052
	Ours	<b>0.071</b>	<b>0.048</b>	<b>0.073</b>	<b>0.049</b>	<b>0.101</b>	<b>0.048</b>	<b>0.085</b>	<b>0.048</b>
Impulse	PPO	0.151	0.101	0.162	0.103	0.216	0.101	0.178	0.101
	RMA	0.096	0.059	0.098	0.060	0.137	0.060	0.120	0.059
	HIM	0.117	0.080	0.120	0.081	0.186	0.080	0.157	0.080
	Ours (w/ HIM)	0.091	0.061	0.093	0.061	0.139	0.061	0.117	0.061
	Ours	<b>0.086</b>	<b>0.058</b>	<b>0.087</b>	<b>0.059</b>	<b>0.115</b>	<b>0.059</b>	<b>0.099</b>	<b>0.058</b>



(a) Varying short-history length.



(b) Varying long-history length.

Fig. 4: Ablation study on the feedback structure with different history lengths.

incorporating short- and long-horizon input/output histories improves performance. Next, we introduce noise into the predictive model’s input and output to assess its robustness and the impact of prediction accuracy on control. Finally, we compare our method using the predictive model versus a full-dynamics model to evaluate the effect of model simplification. All experiments involve quadrotors tracking zigzag trajectories across three random seeds.

1) *History structure*: We evaluate feedback structures incorporating historical records. In short-history experiments, we concatenate recent input/output pairs, including predictive errors, into a single input vector. For long-history experiments, we retain the latest state-action pair while encoding additional history using a TCN model [12], [14], concatenating its latent output with the policy input.

As shown in Fig. 4, using longer history—either short-

TABLE IV: Evaluation of the internal predictive model’s noise tolerance, tested on quadrotors tracking zigzag trajectories.

	Noise $\sigma_n$	Train		Eval	
		static	disturbed	static	disturbed
	0.	0.057	0.064 (+12.28%)	0.101	0.119 (+17.82%)
input noise	0.02	0.063	0.069 (+9.52%)	0.092	0.117 (+27.17%)
	0.04	0.070	0.094 (+34.28%)	0.100	0.154 (+54.00%)
	0.08	0.071	0.121 (+70.42%)	0.134	0.202 (+50.74%)
	0.16	0.082	0.135 (+64.63%)	0.152	0.231 (+51.97%)
output noise	0.02	0.063	0.071 (+12.69%)	0.076	0.132 (+73.68%)
	0.04	0.073	0.097 (+32.87%)	0.102	0.161 (+57.84%)
	0.08	0.074	0.121 (+63.51%)	0.118	0.179 (+51.69%)
	0.16	0.076	0.130 (+71.05%)	0.138	0.222 (+60.86%)

or long-term—does not improve performance over immediate feedback (i.e., the latest one-step history). While differing from [14], this aligns with  $\mathcal{L}_1$  adaptive control [21] and robust control methods [2], [3], where immediate error feedback suffices for accuracy and robustness. This simpler structure also reduces implementation complexity and mitigates cumulative errors in real robots.

2) *Noise tolerance*: To assess the impact of the internal predictive model’s accuracy on control performance, we introduce varying levels of noise into the model’s state input or predictions. Specifically, we introduce zero-mean Gaussian noise with varying variances to the state vector, including position, velocity, quaternion, and angular velocity. For the quaternion, we treat the sampled Gaussian noise as a delta Euler angle, convert it to a delta quaternion, and apply it to rotate the original orientation. We train and evaluate multiple policies using a range of noisy predictive models for quadrotor tracking on zigzag trajectories, with the results presented in Tab. IV.

The results show that our method maintains competitive performance under mild noise, whether applied to model inputs or responses. However, as noise levels increase, the accuracy of the predictive model becomes increasingly crucial, particularly in scenarios with external disturbances or out-of-distribution system dynamics. Across different training and evaluation ranges, performance declines as noise intensifies, with a similar trend observed under external force disturbances. This observation highlights an advantage of our explicit predictive model over other learning-based approaches: it provides

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 5: Real-world experiments on a quadrotor with two different rope-suspending payloads.

TABLE V: Evaluation of the internal predictive model’s complexity, tested on quadrotors tracking zigzag trajectories.

Model	Train		Eval	
	static	disturbed	static	disturbed
Simplified	0.057	0.064 (+12.28%)	0.101	0.119 (+17.82%)
Full	0.061	0.067 (+9.84%)	0.095	0.105 (+10.53%)

a stable baseline estimation, albeit approximate. Although modeling noise can affect feedback reliability, the rigid-body transformation adheres to fundamental motion principles and generalizes across robotic systems. While this model simplifies nonlinear complexities, the RL part in our framework compensates effectively within an informative closed-loop system.

3) *Model simplification*: We investigate the impact of predictive model simplification on performance. A quadrotor serves as our testbed due to its well-studied dynamics model [40]. We compare our approach with a variant that uses the full-dynamics model for tracking zigzag trajectories. As shown in Tab. V, the results indicate that despite the simplification, performance degradation is marginal. Notably, this simplification improves generalization across different morphologies, making it applicable to complex systems that are difficult to model.

### E. Real-world Deployment

To further validate the effectiveness of our framework, we conducted a real-world experiment on a quadrotor, deploying our method to assess its robustness under various disturbances. The quadrotor, equipped with an Nvidia Jetson Xavier NX [42] and a PX4 flight control unit (FCU) [43], interfaces with a motion capture system for real-time pose tracking. The whole neural control framework runs on the embedded computer. The control policy generates CTBR commands, transmitted to the FCU via MAVROS [44] and converted to motor velocities. The quadrotor weighs approximately 1 kg, and we evaluated its tracking performance across five trials under three disturbance conditions: no disturbance, a 200g payload, and a 330g payload, each attached via a rope to increase instability and nonlinearity. The deployment system setup is shown in Fig. 5a.

We present the tracking performance for each axis, along with the normalized error across all axes, in Fig. 5b. The lines represent the average tracking errors under different conditions, with black intervals indicating variance. Our results show that the framework maintains relatively stable performance across various payloads. To illustrate the impact of predictive error feedback, we include a detailed example trajectory for the 330g payload as an ablation study in a real-world scenario (Fig. 5c). The z-axis trajectory demonstrates that, without predictive error feedback, the control policy struggles to maintain a stable flight height due to the unexpected heavy payload. In contrast, our method, which integrates predictive error feedback, achieves robust tracking despite the downward force from the additional payload.

## V. CONCLUSION

In this paper, we propose Neural Internal Model Control (NeuralIMC), a framework that bridges model-based control with RL-based neural control. By leveraging body dynamics and a white-box rigid-body transformation, we avoid high-dimensional nonlinearity, using a simplified predictive model in a closed-loop feedback structure. Model-free RL enables the policy to react to predictive errors. NeuralIMC outperforms state-of-the-art controllers on quadrotors and quadrupedal robots. Real-world tests on a heavily loaded quadrotor demonstrate robustness to disturbances and dynamic variations, highlighting its potential for interpretable and adaptive learning-based control. For the future work, we’d like to conduct more expanded real-world experiments including testing on diverse robots and improving disturbance training [45].

## REFERENCES

- [1] B. Pasik-Duncan, “Adaptive control [second edition, by karl j. astrom and bjorn wittenmark, addison wesley (1995)],” *IEEE Control Systems Magazine*, vol. 16, no. 2, pp. 87–, 1996.
- [2] N. T. Nguyen and N. T. Nguyen, *Model-reference adaptive control*. Springer, 2018.
- [3] D. E. Rivera, M. Morari, and S. Skogestad, “Internal model control: Pid controller design,” *Industrial & engineering chemistry process design and development*, vol. 25, no. 1, pp. 252–265, 1986.
- [4] K. J. Åström and B. Wittenmark, “On self tuning regulators,” *Automatica*, vol. 9, no. 2, pp. 185–199, 1973.

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

- [5] F. Lin, R. D. Brandt, and G. Saikalis, "Self-tuning of pid controllers by adaptive interaction," in *Proceedings of the 2000 American Control Conference. ACC (IEEE Cat. No. 00CH36334)*, vol. 5. IEEE, 2000, pp. 3676–3681.
- [6] B. A. Francis and W. M. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, no. 5, pp. 457–465, 1976.
- [7] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne, "Deepmimic: Example-guided deep reinforcement learning of physics-based character skills," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 143:1–143:14, Jul. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3197517.3201311>
- [8] Y. Li, J. Li, W. Fu, and Y. Wu, "Learning agile bipedal motions on a quadrupedal robot," *arXiv preprint arXiv:2311.05818*, 2023.
- [9] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.
- [10] Y. Song, A. Romero, M. Müller, V. Koltun, and D. Scaramuzza, "Reaching the limit in autonomous racing: Optimal control versus reinforcement learning," *Science Robotics*, vol. 8, no. 82, p. eadg1462, 2023.
- [11] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaa5872, 2019. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.aau5872>
- [12] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *arXiv preprint arXiv:2107.04034*, 2021.
- [13] J. Long, Z. Wang, Q. Li, L. Cao, J. Gao, and J. Pang, "Hybrid internal model: Learning agile legged locomotion with simulated robot response," in *The Twelfth International Conference on Learning Representations*, 2024.
- [14] Z. Li, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, "Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control," *arXiv preprint arXiv:2401.16889*, 2024.
- [15] I. Radosavovic, T. Xiao, B. Zhang, T. Darrell, J. Malik, and K. Sreenath, "Real-world humanoid locomotion with reinforcement learning," *Science Robotics*, vol. 9, no. 89, p. eadi9579, 2024.
- [16] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," *IEEE International Conference on Robotics and Automation*, 2017.
- [17] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *5th Annual Conference on Robot Learning*, 2021. [Online]. Available: <https://openreview.net/forum?id=wK2fDDJ5VcF>
- [18] M. O'Connell, G. Shi, X. Shi, K. Aizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung, "Neural-fly enables rapid learning for agile flight in strong winds," *Science Robotics*, vol. 7, no. 66, p. eabm6597, 2022.
- [19] Z. Li, J. Zeng, A. Thirugnanam, and K. Sreenath, "Bridging model-based safety and model-free reinforcement learning through system identification of low dimensional linear models," *arXiv preprint arXiv:2205.05787*, 2022.
- [20] K. Huang, R. Rana, A. Spitzer, G. Shi, and B. Boots, "DATT: Deep adaptive trajectory tracking for quadrotor control," in *7th Annual Conference on Robot Learning*, 2023. [Online]. Available: <https://openreview.net/forum?id=XEW-cnNsr6>
- [21] F. L. Lewis, "L1 adaptive control theory: Guaranteed robustness with fast adaptation (hovakimyan, n. and cao, c.; 2010 [bookshelf]," *IEEE Control Systems Magazine*, vol. 31, no. 5, pp. 112–114, 2011.
- [22] S. Mallikarjunan, B. Nesbitt, E. Kharisov, E. Xargay, N. Hovakimyan, and C. Cao, "L1 adaptive controller for attitude control of multirotors," in *AIAA guidance, navigation, and control conference*, 2012, p. 4831.
- [23] J. Pravitra, K. A. Ackerman, C. Cao, N. Hovakimyan, and E. A. Theodorou, "L1-adaptive mppi architecture for robust and agile control of multirotors," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 7661–7666.
- [24] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandlekar, B. Babich, G. State, M. Hutter, and A. Garg, "Orbit: A unified simulation framework for interactive robot learning environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3740–3747, 2023.
- [25] B. Xu, F. Gao, C. Yu, R. Zhang, Y. Wu, and Y. Wang, "Omnidrones: An efficient and flexible platform for reinforcement learning in drone control," *IEEE Robotics and Automation Letters*, 2023.
- [26] D. Zhang, A. Loquercio, X. Wu, A. Kumar, J. Malik, and M. W. Mueller, "Learning a single near-hover position controller for vastly different quadcopters," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1263–1269.
- [27] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [28] I. Koryakovskiy, M. Kudruss, H. Vallery, R. Babuška, and W. Caarls, "Model-plant mismatch compensation using reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2471–2477, 2018.
- [29] A. Kumar, Z. Li, J. Zeng, D. Pathak, K. Sreenath, and J. Malik, "Adapting rapid motor adaptation for bipedal robots," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1161–1168.
- [30] M. A. Henson and D. E. Seborg, "An internal model control strategy for nonlinear systems," *AIChE Journal*, vol. 37, no. 7, pp. 1065–1081, 1991.
- [31] M. Bin, P. Bernard, and L. Marconi, "Approximate nonlinear regulation via identification-based adaptive internal models," *IEEE Transactions on Automatic Control*, vol. 66, no. 8, pp. 3534–3549, 2020.
- [32] 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.
- [33] L. Von Stumberg, V. Usenko, and D. Cremers, "Direct sparse visual-inertial odometry using dynamic marginalization," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2510–2517.
- [34] Intel RealSense, "Visual-inertial tracking case study," 2024. [Online]. Available: <https://www.intelrealsense.com/visual-inertial-tracking-case-study/>
- [35] Z. Teed, L. Lipson, and J. Deng, "Deep patch visual odometry," in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 39033–39051. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/7ac484b0f1a1719ad5be9aa8c8455fbb-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/7ac484b0f1a1719ad5be9aa8c8455fbb-Paper-Conference.pdf)
- [36] I. Newton, *Philosophiæ Naturalis Principia Mathematica*. London: Royal Society, 1687.
- [37] F. S. Grassia, "Practical parameterization of rotations using the exponential map," *Journal of graphics tools*, vol. 3, no. 3, pp. 29–48, 1998.
- [38] PX4, *Controller Diagrams — PX4 User Guide*, 2024. [Online]. Available: [https://docs.px4.io/main/en/flight\\_stack/controller\\_diagrams.html](https://docs.px4.io/main/en/flight_stack/controller_diagrams.html)
- [39] E. Kaufmann, L. Bauersfeld, and D. Scaramuzza, "A benchmark comparison of learned control policies for agile quadrotor flight," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 10504–10510.
- [40] M. Faessler, D. Falanga, and D. Scaramuzza, "Thrust mixing, saturation, and body-rate control for accurate aggressive quadrotor flight," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 476–482, 2016.
- [41] E. Süli and D. F. Mayers, *An introduction to numerical analysis*. Cambridge university press, 2003.
- [42] NVIDIA Corporation, "Jetson xavier nx developer kit," <https://developer.nvidia.com/embedded/jetson-xavier-nx>.
- [43] L. Meier et al., "Px4: A professional autopilot," <https://px4.io> and <https://github.com/PX4/PX4-Autopilot>, 2021.
- [44] MAVROS contributors, "Mavros: Mavlink extendable communication node for ros," <https://github.com/mavlink/mavros>.
- [45] J. Long, W. Yu, Q. Li, Z. Wang, D. Lin, and J. Pang, "Learning h-infinity locomotion control," in *8th Annual Conference on Robot Learning*, 2024.