

Global End-Effector Pose Control of an Underactuated Aerial Manipulator via Reinforcement Learning

Shlok Deshmukh, Javier Alonso-Mora, and Sihao Sun*

Abstract—Aerial manipulators, which combine robotic arms with multi-rotor drones, face strict constraints on arm weight and mechanical complexity. In this work, we study a lightweight 2-degree-of-freedom (DoF) arm mounted on a quadrotor via a differential mechanism, capable of full six-DoF end-effector pose control. While the minimal design enables simplicity and reduced payload, it also introduces challenges such as underactuation and sensitivity to external disturbances. To address these, we employ reinforcement learning, training a Proximal Policy Optimization (PPO) agent in simulation to generate feedforward commands for quadrotor acceleration and body rates, along with joint angle targets. These commands are tracked by an incremental nonlinear dynamic inversion (INDI) attitude controller and a PID joint controller, respectively. Flight experiments demonstrate centimeter-level position accuracy and degree-level orientation precision, with robust performance under external force disturbances—including manipulation of heavy loads and pushing tasks. The results highlight the potential of learning-based control strategies for enabling contact-rich aerial manipulation using simple, lightweight platforms. Videos of the experiment and the method are summarized in <https://youtu.be/bWLTpQKcCOA>.

I. INTRODUCTION

Aerial manipulation combines the dexterity of robotic arms with the mobility of quadrotors, enabling in-flight tasks such as pushing, sliding, grasping, and transporting objects [1], [2], [3]. These capabilities are valuable for applications including material transport, disaster response, industrial inspection and maintenance, where they can reduce cost and risk for human workers. However, deploying aerial manipulators in the real world is challenging due to the design constraints (space, weight, and energy), as well as control challenges due to nonlinear coupled dynamics between the flying base and the arm, together with disturbances arising from interaction with the environment.

Many aerial manipulators employ redundant arms with multiple degrees of freedom, or add additional tilted / tiltable rotors to achieve full actuation of the base, but at the cost of additional weight and structural complexity (e.g., [4], [5], [6]). Unlike ground and underwater robots, aerial robots must prioritize lightweight design due to their limited thrust-to-weight ratio and restricted onboard energy resources. To address this, we adopt a minimal yet effective design: the *Differential Shoulder Aerial Manipulator (DSAM)*, where a single rigid arm with an end-effector (e.g., a gripper) is connected to the quadrotor base through a differential driven by two rotational actuators, enabling pitch and roll




Fig. 1: Snapshot of a DSAM aerial manipulator, consisting of a quadrotor base with an arm mounted through a differential mechanism that provides 2-DoF motion.

motions (Figure 1). Prior work has shown that a configuration for AM like DSAM is sufficient to achieve *global* 6-DoF end-effector pose regulation, despite its mechanical simplicity [7]. However, the problem of designing robust whole-body control algorithms for DSAM remains largely unexplored.

Before developing model-based control methods, we first investigate reinforcement learning (RL) for the whole-body control of DSAM. RL-based controllers offer significant advantages by alleviating the need for analytical models and their simplifications, implicitly considering the system actuation constraints through explorations, enabling fast inference, and supporting adaptation to unmodeled disturbances through domain randomization, all of which are essential for real-world aerial manipulation. To address the sim-to-real gap, we also employ domain randomization during training, and adopt the scheme that hybridizes the RL with downstream low-level controllers [8]. Specifically, we employ an incremental nonlinear dynamic inversion (INDI) attitude controller for the quadrotor base, and a PID controller for the arm. We have experimentally validated the proposed algorithm in controlling the 6-DoF end-effector pose, as well as experiments testing the robustness against external forces on the end-effector, including carrying an additional load, and pushing an object.

Our contributions are summarized as follows:

- 1) We develop and deploy a reinforcement learning policy, trained using Proximal Policy Optimization (PPO), for whole-body control of the Differential Shoulder Aerial Manipulator (DSAM). The policy achieves

Authors are with the Department of Cognitive Robotics (CoR), Faculty of Mechanical Engineering, Delft University of Technology, Delft, Netherlands. * Corresponding author (ssun1@tudelft.nl).

centimeter-level position accuracy and degree-level orientation precision in real-world experiments.

- 2) We demonstrate robust end-effector pose control under significant external disturbances, including lifting a payload of 140 g (over 16% of system mass) and pushing a 590 g object (over 68% of system mass).
- 3) Through ablation studies, we show that body-rate and joint observations, high-level control abstractions, and domain randomization of mass and friction significantly improve sim-to-real transfer, as evidenced by learning curve analysis.

II. RELATED WORK

A. System Design

Aerial manipulators can be grouped into two design families, each aiming to realize end-effector (EE) regulation in SE(3) or a relevant subspace. (i) The first augments an underactuated unidirectional multirotor base with a serial or parallel arm (e.g., [4], [9]). (ii) The second employs overactuated drone bases with tilted or tiltable rotors, allowing EEs fixed to the base to realize arbitrary pose tracking in SE(3) (e.g., [10], [11], [12], [13], [14]). Some designs employ both choices (e.g., [5], [6], [15]). A more detailed testimony of aerial manipulator designs, and their comparisons, is given in the survey [1]. We focus on the first family and adopt the DSAM: a quadrotor base connected to an arm through a 2-DoF differential. DSAM is one of the most minimal designs achieving *global* 6-DoF end-effector pose control, where the dual rotational actuators enlarge arm torques, thereby boosting load and manipulation capability without increasing weight or complexity. Ref. [7] concludes that task flatness holds for a DSAM in a special condition, namely, the differential is located at the CoM of the quadrotor base. In practice, however, this condition may not hold due to design constraints and CoM variations. In addition, directly designing the control law through differential flatness cannot effectively satisfy input and state constraints, which are essential for aerial robotic manipulators.

B. Planning and Control

Specifically for DSAM, a controller still needs to be designed to reliably track setpoints and trajectories across the full SE(3) task space. Most research on aerial manipulator control employ model-based methods. Reactive control methods, such as feedback linearization and geometric control [16], [17], [18], [19], [20], offer computational efficiency but provide limited mechanisms for handling actuator bounds and state constraints. Several works employ model predictive control (MPC) and Model Predictive Path Integral (MPPI) control for aerial manipulators, which improve constraint handling and can reason over short horizons, yet they incur substantial online computation and require careful contact and cost-shaping design in contact-rich tasks [21], [22], [23]. Motivated by promising progress in robot learning for manipulation (e.g., [24]), recent studies have begun to explore reinforcement learning for aerial manipulators,

demonstrating promising results on platforms with fully-actuated bases (e.g., omnidirectional bases with fixed EEs for door opening [25]). Building on this trend, we present the first whole-body controller for DSAM, based on learning-based method, that globally tracks SE(3) setpoints and addresses the underactuation of the DSAM while embracing its hardware design simplicity. This work will pave the way to future research for DSAM in contact-rich aerial manipulation tasks.

III. PRELIMINARIES

Three coordinate frames are considered in this work:

- 1) World frame $\{x_W, y_W, z_W\}$ - inertial frame of reference.
- 2) Body frame $\{x_B, y_B, z_B\}$ - located at the drone platform's CoM.
- 3) End-effector frame $\{x_E, y_E, z_E\}$ - located at gripper's CoM.

Note that for dynamic consistency, we define the end-effector frame at the CoM of the gripper, rather than at its tip. Modeling each link and arm as a rigid body, we define the configuration and velocity of the system as

$$\mathbf{q} = \begin{Bmatrix} {}^W \mathbf{p}_b \\ \mathbf{R}_{wb} \\ \boldsymbol{\theta} \end{Bmatrix} \in \mathbb{R}^3 \times SO(3) \times \mathbb{T}^2, \quad \mathbf{v} = \begin{Bmatrix} {}^W \dot{\mathbf{p}}_b \\ {}^B \boldsymbol{\Omega}_b \\ \dot{\boldsymbol{\theta}} \end{Bmatrix} \in \mathbb{R}^8, \quad (1)$$

where ${}^W \mathbf{p}_b$ is the position of the system in world frame, \mathbf{R}_{wb} the rotation from world to body frame and $\boldsymbol{\theta}$ are joint angles.

The system evolves according to coupled nonlinear dynamics given by generalized Lagrangian method presented in [26]

$$\mathbf{M}(\mathbf{q})\dot{\mathbf{v}} + \mathbf{C}(\mathbf{q}, \mathbf{v})\mathbf{v} + \mathbf{G}(\mathbf{q}) = \boldsymbol{\tau}(\mathbf{q}, \mathbf{u}), \quad (2)$$

where $\boldsymbol{\tau}$ represents the generalized forces due to control inputs $\mathbf{u} := [\boldsymbol{\omega}_{cmd}, \boldsymbol{\tau}_\theta]^\top$ and configuration \mathbf{q} . Here, $\boldsymbol{\omega}_{cmd} \in \mathbb{R}_{\geq 0}^4$ denotes rotor speed commands for the quadrotor base, and $\boldsymbol{\tau}_\theta \in \mathbb{R}^2$ denotes joint torques.

Our objective is to control the end-effector pose of the system, defined by the forward kinematics:

$${}^W \mathbf{x}_e := [{}^W \mathbf{p}_e, \mathbf{q}_{we}]^\top = f_f(\mathbf{q}) \in SE(3) \quad (3)$$

where ${}^W \mathbf{p}_e$ is end-effector position in world frame and \mathbf{q}_{we} is the quaternion from the world to end-effector frame.

We formulate the control task learning a policy at each time t , maps the observed state $s_t \in \mathcal{S}$ to an action $a_t \in \mathcal{A} \subseteq \mathbb{R}^9$, which is then mapped to valid control inputs $\mathbf{u}_t \in \mathbb{R}^6$ using low-level controllers (*LLC*):

$$\pi_\phi : \mathcal{S} \rightarrow \mathcal{A}, \quad a_t = \pi_\phi(s_t), \quad \mathbf{u}_t = LLC(a_t) \quad (4)$$

During learning, at time t the task is scored by an instantaneous reward:

$$r_t = r(s_t, a_t) \quad (5)$$

The control inputs \mathbf{u}_t directly influence the system dynamics through (2). The rewards aim to minimize the tracking error between the actual and desired end-effector poses

consumer grade GPU (NVIDIA RTX 2080 Ti).

2) *Rewards*: We define two primary rewards, one for position and one for orientation, along with additional penalty terms to smooth the action space and penalize large action magnitudes. All rewards are scaled by a negative exponential transformation that maps raw reward values between $(0, 1]$, enabling comparable weighing of individual reward components.

Position reward: We use the L2 norm of the difference between the end-effector and goal positions as the basis for position reward:

$$r_{\text{pos}} = w_1 \cdot \exp(-\alpha_1 \cdot \|{}^W \mathbf{p}_e - {}^W \mathbf{p}_{\text{goal}}\|_2). \quad (6)$$

The α_1 value tunes reward sensitivity for local convergence. Larger values increase the gradient magnitude near the target but can lead to vanishing gradients far from the target, as the exponential saturates near zero.

Orientation reward: We use the smallest geodesic angular difference between the orientations of the end-effector and the goal $\mathcal{E}_{\text{ori}} = 2 \arccos(|\mathbf{q}_e \cdot \mathbf{q}_{\text{goal}}|)$ to represent the orientation error of the end-effector. We do not penalize the quadrotor’s orientation directly, as the end-effector orientation already constrains the possible quadrotor orientations:

$$r_{\text{ori}} = w_2 \cdot \exp(-\alpha_2 \cdot \mathcal{E}_{\text{ori}}). \quad (7)$$

Quadrotor action smoothing reward: The norm of the difference between the quadrotor base action values ${}^{\text{base}} \mathbf{a} = \{{}^W \mathbf{a}_b, {}^B \boldsymbol{\Omega}_b, {}^W \psi_r\}$ between the current and previous time steps penalizes large deviations in consecutive actions:

$$r_{\text{ds}} = w_3 \cdot \exp(-\alpha_3 \cdot \|{}^{\text{base}} \mathbf{a}_{t-1} - {}^{\text{base}} \mathbf{a}_t\|_2^2). \quad (8)$$

We found this term beneficial for reliable deployment.

Joint action smoothing reward: The absolute difference between joint position action commands (${}^{\text{joint}} \mathbf{a} = \boldsymbol{\theta}_r$) is penalized to encourage smoother joint actions, producing near-constant torque with fewer jerks that might destabilize the platform:

$$r_{\text{js}} = w_4 \cdot \exp(-\alpha_4 \cdot \|{}^{\text{joint}} \mathbf{a}_{t-1} - {}^{\text{joint}} \mathbf{a}_t\|_1). \quad (9)$$

Quadrotor action magnitude reward: The magnitude of quadrotor base actions ${}^{\text{base}} \mathbf{a} = \{{}^W \mathbf{a}_b, {}^B \boldsymbol{\Omega}_b, {}^W \psi_r\}$ is penalized to minimize oscillations and overshoot while reaching the goal pose:

$$r_{\text{dmag}} = w_5 \cdot \exp(-\alpha_5 \cdot \|{}^{\text{base}} \mathbf{a}_t\|_2^2). \quad (10)$$

The total reward is the weighted sum of all components:

$$\mathcal{R}_t = r_{\text{pos}} + r_{\text{ori}} + r_{\text{ds}} + r_{\text{js}} + r_{\text{dmag}}. \quad (11)$$

3) *References and Resets*: We sample goal pose references¹ for the end-effector (gripper) during training. A new goal pose is sampled upon environment resets, which occur when an episode terminates either due to a timeout or when the quadrotor falls below 0.3 m.

¹We use the terms *references* and *commands* interchangeably, following Isaac Lab terminology for references in goal-conditioned tasks

Index	Weight w_i	Scale α_i
1	4.0	1.2
2	1.0	1.0
3	0.5	1.0
4	1.0	1.0
5	0.1	1.0

TABLE I: Reward weights w_i and scaling factors α_i used in the total reward function.

Positions are sampled from the ranges $(-1.0, 1.0)$ meter in the X - Y directions and $(3.0, 5.0)$ m in the Z direction. Orientations are sampled from $(-180^\circ, +180^\circ)$ in yaw and $(-90^\circ, +90^\circ)$ in both pitch and roll, covering the reachable workspace.

After resets, the quadrotor base spawns at the same initial position $(0.0, 0.0, 3.0)$ meters. However, the manipulator configuration is randomly sampled from $(-90^\circ, +90^\circ)$ for each joint, which encourages the quadrotor to learn to track poses from arbitrary arm configurations.

4) *Domain Randomization*: After every reset, we randomly sample a payload mass to be added to the end-effector from $[-15\text{g}, 120\text{g}]$ and scale inertia accordingly. This results in a generalized policy that is able to perform pose tracking while carrying loads up to 140 g in the real world, shown in Section V-C.

We also observe noticeable differences in joint friction between the training and deployment environments, leading to variations in joint behavior under the same control inputs. To address this, we first improve our simulation fidelity by finding friction and stiffness values that result in similar step responses from joint actuators across training, deployment simulators and the hardware. Once this gap is reduced, we apply domain randomization by scaling joint stiffness and friction to $[0.75, 1.25]$ and $[0.0, 1.5]$ times their default values, respectively. The necessity of domain randomization is also analyzed through ablation studies in Section VI-C.

V. REAL-WORLD EXPERIMENTS

A. Experimental Setup

The policy trained using the proposed method is deployed on our 2-DoF DSAM, with a total weight of 860g. The moving part (arm plus gripper) weighs 75 g, actuated by two Dynamixel-XC330 motors [34] with embedded PID controllers following joint commands. Our control architecture is developed upon Agilicious [35], a modular software framework designed for agile quadrotor flight. The control stack entirely runs on the DSAM’s onboard CPU (Raspberry Pi 5), including onboard inference using ONNX Runtime [27]. This results in an average inference time of **0.18 ms**.

Experiments are conducted in a 13×7 m facility equipped with a motion capture system that provides ground-truth pose measurements at 100 Hz, which is then fused with an EKF running onboard at 300 Hz to obtain the state estimates. A safety controller that only controls the quadrotor base motion [29] is used for take-off and handles workspace violations by immediately bringing the system to hover state.

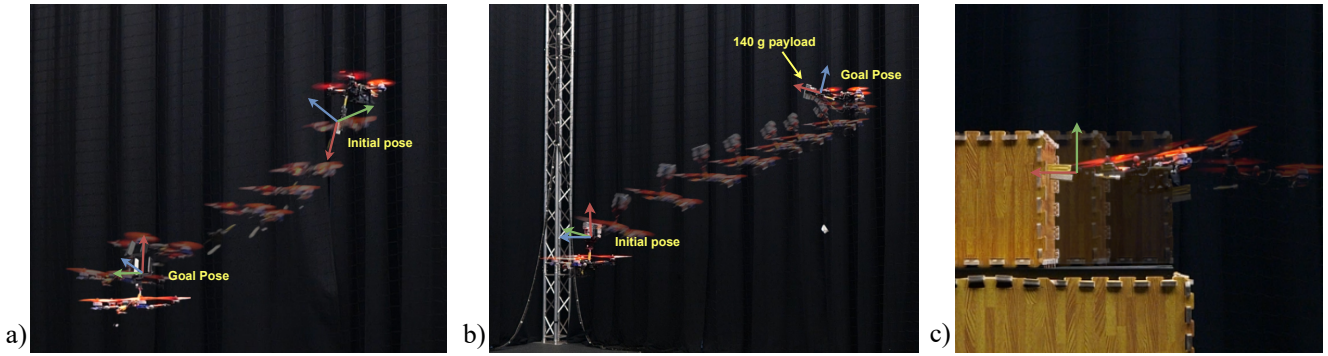


Fig. 3: a) End-effector pose control between two randomly sampled goal poses, b) Pose control with a 140 g payload held by the gripper, c) DSAM pushing a box weighing 590 g from right to left.

After take-off, control is switched from the safety controller to the learned policy that actuates the quadrotor to reach the commanded end-effector pose.

B. End-Effector Pose Control

We evaluate end-effector pose control on physical system by sending 10 randomly sampled goal poses from a bounded workspace: $p_{x,goal} \in [-1, 1]$, $p_{y,goal} \in [-1, 1]$, $p_{z,goal} \in [1, 2]$ with roll and pitch angles from -90° to 90° and yaw angle from -120° to 120° . Each pose is maintained for 10 seconds, giving time for convergence, after which the next goal is sent. A pose is considered achieved, if the end-effector converges to the commanded pose without any intervention. The snapshot of the experiment is given in Figure 3a. The tracking performance is summarized in Table II and presented in Figure 4.

C. End-Effector Pose Control with Additional Load

To evaluate robustness under load we perform pose control while carrying payloads of 50 g and 140 g on 7 randomly generated goals. A snapshot of this experiment is presented in Figure 3b.

Condition	Position Error (m)		Orientation Error (deg)		Success Rate (# reached / total)
	mean	std.	mean	std.	
0 g	0.0536	0.0166	8.8078	7.1834	10/10
50 g	0.0995	0.0695	12.5020	3.0619	7/7
140 g	0.0954	0.0505	15.7006	4.8312	7/7

TABLE II: Pose tracking error (mean and standard deviation) and success rate across payload conditions.

Our system maintains an average position error of 10 cm for both 50 g and 140 g payloads. The orientation error increases by 3° with 140 g, which we attribute to the manipulator arm’s higher sensitivity to load induced disturbances than the quadrotor base, visible from the quaternion component responses in Figure 4. Position tracking errors are observed at extreme joint angles (visible from 30 - 40 s in Figure 4).

It is worth noting that with a 140 g payload, the moving part (arm plus load) accounts for 215 g, or 21.5% of the

total DSAM mass (DSAM plus load). Despite this, the policy maintains stable flight and demonstrates consistent convergence under load. The proposed controller also achieves agile transient responses with minimal overshoot, while preserving stability. However, although domain randomization improves position tracking, we argue that it alone is still insufficient for accurate pose tracking under disturbances. A greater pose error has been observed in comparison with the baseline. An improved low-level controller capable of compensating for such disturbances may be required to further enhance tracking performance.

D. Pushing an Object

To evaluate robustness against contact-induced disturbances and the ability to exert forces on the environment, we task the policy with pushing a 590 g box along a straight line while maintaining a fixed end-effector orientation. We generate a straight-line path with constant end-effector orientation and provide the corresponding poses to the policy. We then sample target poses along the path to ensure a constant end-effector velocity and feed them sequentially to the policy. Following these references, the DSAM approaches the box, establishes contact, and successfully pushes it. During contact, the DSAM remains stable and continues to maintain end-effector orientation while adjusting the quadrotor attitude to exert a force, moving the box by 30 cm (Figure 3c). This result demonstrates that the learned whole-body controller can tolerate sudden interaction forces, despite being trained without explicit contact models or force feedback.

E. Path following

As an extension to pose control, we attempt path following without retraining. The end-effector is commanded to track a figure-eight path. The same as the pushing task, we achieve path following by sampling a sequence of target poses along the path and feeding them to the policy one at a time. The 3D trajectory during the path following task is presented in Figure 5.

As seen in Table III, the RMSE for the figure-8 path is noticeably larger than for the straight line, reflecting the higher




Fig. 4: Comparison of pose control performance with and without payload (we show only the first 5 setpoints for readability).

Path	Position RMSE (m)	Orientation RMSE (deg)
Figure-8	0.8167	14.3915
Straight line	0.1960	5.5607

TABLE III: RMSE of position and orientation errors for the two path following tasks.

demands of multi-axis coordination compared to single-axis motion. The tracking accuracy is further limited because the policy is conditioned only on the immediate target pose and does not explicitly account for higher-order derivatives such as velocity, acceleration, or jerk. Consequently, the policy was not trained specifically for accurate path tracking but rather for pose regulation. Future work could address this by augmenting the observation space with future waypoints or target velocities, which have been shown to improve path-tracking performance in related learning-based control frameworks (e.g., [36]). Despite these limitations, the policy is able to follow simple paths while maintaining orientation, demonstrating whole-body coordination emerging from pose control alone.

VI. ABLATION STUDIES

To study the impact of design choices, we perform ablations on various components of our training setup, such as the observation space, low-level controller choice, domain ran-

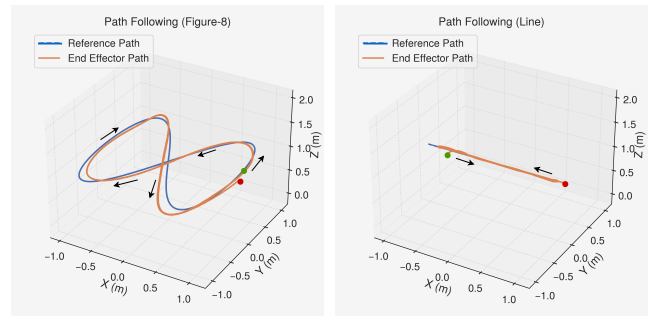


Fig. 5: 3D trajectory of path followed by end-effector (starting point indicated by green dot)

domization, and reward function. We summarize the results in the learning curves, as well as position error distributions, presented in Figure 6.

A. Observation Space Ablations

1) *Joint positions*: For a 2D-DSAM, through inverse kinematics, a unique mapping exists from the end-effector pose error and the goal position in the quadrotor frame (current observation) to the joint positions. In other words, having joint positions leads to redundancies in the observation space. However, despite being theoretically redundant to have joint positions in the observation, we notice that including the joint positions in the observation space leads to a significant increase in end-effector orientation reward.

2) *Goal positions in quadrotor frame*: In the observation space, we include the quadrotor position information by expressing the goal position in the quadrotor frame. Although these positions could, in principle, be derived from the goal pose in the end-effector frame and the joint angles, which have been included in the observation space, via inverse kinematics, omitting them leads to a substantial drop in position rewards.

A plausible explanation is that explicitly providing the joint positions and goal position in the quadrotor frame simplifies the learning problem by bypassing the need for the policy to infer this relationship implicitly, thereby improving training efficiency and tracking performance.

3) *Body rate of the quadrotor base*: We noticed that including body rate in the observation space is also critical for training for DSAM, which is particularly different from previous works on quadrotors [37], and AMs with a fully-actuated base [25]. We argue that including quadrotor angular rates in the observation significantly improves policy training, as they allow the policy to indirectly infer the effect of torque disturbances induced by the arm. This additional information helps the policy learn the coupled base–arm dynamics especially for DSAM with an underactuated unidirectional base, which is, on the other hand, less critical for omnidirectional platforms or simple quadrotors.

B. Inner-Loop Controller Ablations

While incorporating an inner-loop is a common practice in RL-based control to achieve reliable sim-to-real transfer,




Fig. 6: a) Learning curves with ablations in the training setup, b) Position error distribution on the end-effector pose control task carrying a 100 g load, without mass randomization, c) Oscillating joint position command from policy trained without friction randomization.

determining the appropriate level of control abstraction is non-trivial. In particular, for the quadrotor base of DSAM, we replace the attitude controller with a simple proportional controller that maps the body-rate command to the desired body angular acceleration, which is then passed to the low-level INDI controller. Under this setup, we evaluate the traditional “collective-thrust plus body rate” (CTBR) output of the RL policy for quadrotor flight, following prior work [37]. INDI is retained, as it plays a critical role in handling disturbances, as highlighted in previous studies [29], [38]. The training curves in Fig. 6 show that our setup, where the attitude controller accepts reference acceleration (or equivalently attitude) from the RL policy, achieves consistently higher rewards in both orientation and position tracking.

C. Domain Randomization Ablations

1) *End-effector mass*: As shown in Fig. 6b, mass randomization leads to higher position accuracy while carrying additional load (100g), which can be attributed to the additional exploration induced by varying inertia. When exposed to different mass conditions, the agent is compelled to develop more robust control actions. In this sense, mass randomization serves as a form of controlled exploration, preventing the policy from overfitting to a single configuration and ultimately improving position-tracking performance.

2) *Joint friction*: By removing joint friction randomization, we observe a substantial drop in orientation learning rewards. Similar to the previous ablation, the setup with constant friction overfits to a narrow set of end-effector mass and joint stiffness values. When these properties are randomized with constant friction, a stick and slip motion is observed, resulting in unpredictable arm dynamics and reduced orientation learning. Without friction randomization, we observe oscillations in the joint actions generated by the policy (Figure 6c).

VII. CONCLUSION

We presented an RL-based control method for whole-body control of the DSAM, a minimal 2-DoF aerial manipulator capable of global 6-DoF end-effector pose control. By combining a PPO-based policy with INDI and PID low-level controllers and executing fully onboard with 0.18 ms of inference time, we achieved centimeter-level position

accuracy and degree-level orientation accuracy in real-world experiments, while maintaining stable flight under disturbances, namely, carrying a payload on the end effector (over 16% of the system mass) and pushing an object over 68% of the system mass.

Our study suggests several future directions, including incorporating future waypoints or higher-order reference terms into the observation space to improve trajectory tracking, and extending the framework to contact-rich manipulation tasks. Overall, this work establishes the feasibility of learning-based whole-body control for DSAM, demonstrates its capability in real-world aerial manipulation, and lays the groundwork for increasingly accurate and versatile aerial manipulators.

VIII. ACKNOWLEDGEMENTS

This work was supported by the Dutch Research Council (NWO) through the Veni talent Programme (grant no. 20256, Accurate Aerial Manipulation). We thank Jack Zeng for insightful discussions and for sharing low-level controller templates. We also thank Maurits Pfaff and Kseniia Khomenko for building the 2-DoF DSAM and for their support during experiments.

REFERENCES

- [1] A. Ollero, M. Tognon, A. Suarez, D. Lee, and A. Franchi, “Past, present, and future of aerial robotic manipulators,” *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 626–645, 2022.
- [2] A. Albers, S. Trautmann, T. Howard, T. A. Nguyen, M. Frietsch, and C. Sauter, “Semi-autonomous flying robot for physical interaction with environment,” in *2010 IEEE Conference on Robotics, Automation and Mechatronics*, 2010, pp. 441–446.
- [3] D. Mellinger, Q. Lindsey, M. Shomin, and V. Kumar, “Design, modeling, estimation and control for aerial grasping and manipulation,” in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 2668–2673.
- [4] S. Kim, S. Choi, and H. J. Kim, “Aerial manipulation using a quadrotor with a two dof robotic arm,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2013, pp. 4990–4995.
- [5] M. Tognon et al., “A truly-redundant aerial manipulator system with application to push-and-slide inspection in industrial plants,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1846–1851, 2019.

- [6] A. Afifi, M. van Holland, and A. Franchi, "Toward physical human-robot interaction control with aerial manipulators: Compliance, redundancy resolution, and input limits," in *2022 International conference on robotics and automation (ICRA)*, IEEE, 2022, pp. 4855–4861.
- [7] J. Welde, J. Paulos, and V. Kumar, "Dynamically feasible task space planning for underactuated aerial manipulators," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3232–3239, 2021.
- [8] J. Hwangbo et al., "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, eaau5872, 2019.
- [9] H. Chen, W. Deng, B. Ye, Y. Xiong, Z. Pan, and X. Lyu, "Aerial grasping via maximizing delta-arm workspace utilization," *arXiv preprint arXiv:2506.15539*, 2025.
- [10] R. Rashad, F. Califano, and S. Stramigioli, "Port-hamiltonian passivity-based control on se (3) of a fully actuated uav for aerial physical interaction near-hovering," *IEEE Robotics and automation letters*, vol. 4, no. 4, pp. 4378–4385, 2019.
- [11] J. Li, J. Sugihara, and M. Zhao, "Servo integrated nonlinear model predictive control for overactuated tiltable-quadrotors," *IEEE Robotics and Automation Letters*, 2024.
- [12] M. Kamel et al., "The voliro omniorientational hexacopter: An agile and maneuverable tiltable-rotor aerial vehicle," *IEEE Robotics & Automation Magazine*, vol. 25, no. 4, pp. 34–44, 2018.
- [13] T. Hui, E. Cuniato, M. Pantic, M. Tognon, M. Fumagalli, and R. Siegwart, "Passive aligning physical interaction of fully-actuated aerial vehicles for pushing tasks," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024, pp. 6152–6158.
- [14] M. Tognon and A. Franchi, "Omnidirectional aerial vehicles with unidirectional thrusters: Theory, optimal design, and control," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2277–2282, 2018.
- [15] G. Malczyk, M. Brunner, E. Cuniato, M. Tognon, and R. Siegwart, "Multi-directional interaction force control with an aerial manipulator under external disturbances," *Autonomous Robots*, vol. 47, no. 8, pp. 1325–1343, 2023.
- [16] V. Lippiello and F. Ruggiero, "Exploiting redundancy in cartesian impedance control of uavs equipped with a robotic arm," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2012, pp. 3768–3773.
- [17] G. Zhang, Y. He, B. Dai, F. Gu, J. Han, and G. Liu, "Robust control of an aerial manipulator based on a variable inertia parameters model," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 11, 9515–9525, Nov. 2020, ISSN: 1557-9948.
- [18] Y.-C. Liu and C.-Y. Huang, "Ddpg-based adaptive robust tracking control for aerial manipulators with decoupling approach," *IEEE Transactions on Cybernetics*, vol. 52, no. 8, pp. 8258–8271, 2022.
- [19] V. Lippiello and F. Ruggiero, "Exploiting redundancy in cartesian impedance control of uavs equipped with a robotic arm," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 3768–3773.
- [20] M. Orsag, C. Korpela, M. Pekala, and P. Oh, "Stability control in aerial manipulation," in *2013 American Control Conference*, 2013, pp. 5581–5586.
- [21] G. Garimella and M. Kobilarov, "Towards model-predictive control for aerial pick-and-place," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4692–4697.
- [22] J. Martí-Saumell, J. Sola, A. Santamaria-Navarro, and J. Andrade-Cetto, "Full-body torque-level non-linear model predictive control for aerial manipulation," *arXiv preprint arXiv:2107.03722*, 2021.
- [23] M. Brunner, G. Rizzi, M. Studiger, R. Siegwart, and M. Tognon, "A planning-and-control framework for aerial manipulation of articulated objects," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 689–10 696, 2022.
- [24] A. Rajeswaran et al., "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.
- [25] E. Cuniato, I. Geles, W. Zhang, O. Andersson, M. Tognon, and R. Siegwart, "Learning to open doors with an aerial manipulator," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [26] P. J. From, "An explicit formulation of singularity-free dynamic equations of mechanical systems in lagrangian form—part two: Multibody systems," *Modeling, Identification and Control*, vol. 33, no. 2, pp. 61–68, 2012.
- [27] O. R. developers, *Onnx runtime*, <https://onnxruntime.ai/>, 2021.
- [28] D. Brescianini and R. D'Andrea, "Tilt-prioritized quadcopter attitude control," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 2, pp. 376–387, 2018.
- [29] S. Sun, A. Romero, P. Foehn, E. Kaufmann, and D. Scaramuzza, "A comparative study of nonlinear mpc and differential-flatness-based control for quadrotor agile flight," *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3357–3373, 2022.
- [30] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [31] M. Mittal et al., "Orbit: A unified simulation framework for interactive robot learning environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3740–3747, 2023.
- [32] A. Serrano-Munoz, D. Chrysostomou, S. Bøgh, and N. Arana-Arexolaleiba, "Skr! Modular and flexible library for reinforcement learning," *Journal of Machine Learning Research*, vol. 24, no. 254, pp. 1–9, 2023.
- [33] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, *Learning to walk in minutes using massively parallel deep reinforcement learning*, 2022. arXiv: 2109.11978 [cs.RO].
- [34] ROBOTIS Co., Ltd., *Xc330-m288 dynamixel e-manual*, Accessed: 2025-08-10, 2025. [Online]. Available: <https://emmanual.robotis.com/docs/en/dx1/x/xc330-m288/>.
- [35] P. Foehn et al., "Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight," *Science Robotics*, 2022.
- [36] S. Garg, H. Masnavi, B. Fidan, and F. Janabi-Sharifi, "Dual agent learning based aerial trajectory tracking," *IEEE Robotics and Automation Letters*, 2025.
- [37] 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, eadg1462, 2023.
- [38] R. Ferede, C. De Wagter, D. Izzo, and G. C. De Croon, "End-to-end reinforcement learning for time-optimal quadcopter flight," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024, pp. 6172–6177.