

PINGU: Extending Air-Bearing Spacecraft Emulators with Open-Source Actuators and Learned Control for Contact-Rich Proximity Operations

Ricard Marsal I Castan¹, Akiyoshi Uchida², Aman Arora¹, Pedro Lima¹, Matteo El-Hariry¹, Anrej Orsula¹,
 Francesco Grella¹, Antoine Richard¹, Cedric Pradalier³, Miguel A. Olivarez-Mendez¹
¹University of Luxembourg, ²Tohoku University, ³Georgia Institute of Technology




Fig. 1: **Overview.** Hardware modules (left) and the software stack (right), evaluated on PINGU, our ATMOS-compatible air-bearing platform integration, across the representative use cases shown (bottom).

Abstract—Low-cost planar air-bearing testbeds have matured into a standard proxy for free-flying spacecraft GNC, but they remain largely thruster-only and are rarely equipped for contact-rich, inertia-coupled manipulation. Building on the open-source ATMOS testbed (from which we adopt much of the low-level software and a similar cold-gas thruster configuration) we contribute a reaction wheel and two force/torque-sensed robotic arms (LEVION) with interchangeable end-effectors, integrated as first-class control actuators through a unified ROS 2 abstraction layer. On top of the software stack we build a reinforcement-learning training environment and digital twin, and a controller that exploits these added degrees of freedom, letting classical optimal controllers and learned policies be swapped on the same hardware without modification. We validate the integrated system, PINGU, across four benchmark tasks: point-to-pose navigation (classical LQR vs. sim-to-real PPO), dynamic disturbance rejection under arm-induced center-of-mass shifts, reaction-wheel momentum stabilization, and force-controlled docking. The results show that these additions extend an ATMOS-class emulator

into the contact-rich regime and bridge classical optimal control and reinforcement learning on one reproducible platform.

I. INTRODUCTION

On-orbit servicing, active debris removal, modular satellite assembly, and proximity operations all share a common technical core: vehicles that can simultaneously *translate*, *re-orient*, and *physically interact* with other bodies in micro-gravity. As mission concepts move from rendezvous-and-dock toward dexterous on-orbit manipulation, the dynamics of the spacecraft itself become coupled to the motion of its on-board manipulators: every arm swing redistributes mass, perturbs the center of mass, and induces reaction torques that the attitude-control system must reject. Validating algorithms for this regime on flight hardware is prohibitively expensive, so the

community has long relied on terrestrial *floating platforms* that approximate frictionless planar motion using air bearings on a flat surface. In doing so they offer a low-cost, high-throughput proxy for the planar projection of orbital dynamics, and they have a long, well-surveyed history [1, 2].

a) The limits of single-purpose testbeds: Despite this maturity, the existing landscape of planar emulators is highly fragmented along functional lines. Thruster-only testbeds such as the Luleå and Luxembourg floating platform [3, 4], thruster-propeller testbeds like ATMOS [5], and thruster-and-reaction-wheel testbeds such as the ESA Orbital Robotics Lab platform [6], M-STAR [7], and ASTROS [8], excel at studying guidance, navigation, and control (GNC) for free-flying base motion, but they do not carry on-board manipulators, so the coupled dynamics that arise from arm-induced inertia and CoM shifts cannot be reproduced. Conversely, the much shorter list of testbeds that *do* integrate manipulators - for example York University’s rendezvous-and-capture testbed [9], the FIT Bob & Charlie vehicles [10], and Sabatini et al.’s coordinated-manipulator platform [11] - typically equip a single 3-DoF arm or passive grippers, lack co-integrated thrusters and reaction-wheel control authority, and are demonstrated under classical, hand-tuned controllers on narrow task suites. Astrobees [12, 13], deployed on the ISS and used as a ground emulator, demonstrates the value of a unified flight-and-ground platform, but its single perching arm is not designed for the high-thrust, high-torque, contact-rich experiments that current robotic-servicing research demands. A recent analysis of intra-vehicular free-flyers [14] echoes this conclusion at the system level: in-space manipulation with free-flying robots has been demonstrated only to a limited extent, and the dynamic challenges of a free-floating base coupled to an active manipulator remain a central, under-tested problem. A researcher today who wishes to ask, for example, “*does a learned policy reject CoM shifts induced by arm motion better than an LQR with feed-forward inertia compensation?*” has no single platform on which to run both halves of the comparison cleanly.

b) Our approach: We build directly on the open-source ATMOS testbed [5], adopting much of its low-level software and a similar cold-gas thruster layout on a slightly restructured chassis, and extend it in three parts. First, we add two new control actuators — a high-torque reaction wheel for fine attitude control and momentum management, and two modular arms with 6-axis force/torque sensing and interchangeable end-effectors — exposed alongside the thrusters through a common ROS 2 actuator-abstraction layer, so that any controller (LQR, MPC, or a learned policy) can be swapped without changing the hardware interface. Second, on top of this stack we build a reinforcement-learning training environment and controller, and digital twin based on the RoboRAN [15] and Space Robotics Bench simulator [16], to exploit the additional degrees of freedom the wheel and arms provide. Third, we validate the integrated system in the lab as PINGU, Fig. 1, demonstrating the added actuators and the learned controller across four tasks: point-to-pose navigation, dynamic disturbance rejection under arm-induced CoM shifts, momentum stabilization, and force-controlled docking.

c) Bridging optimal control and reinforcement learning: A second, complementary motivation underlies the platform’s design. Spacecraft GNC has traditionally been the landscape of model-based optimal control, which provides strong theoretical guarantees but assumes a well-characterized rigid-body dynamics model, an assumption that is violated when on-board manipulators move and redistribute the body’s mass. Reinforcement learning (RL) offers a path to controllers that adapt to such structured but hard-to-model disturbances, and recent work has begun to demonstrate sim-to-real RL on free-flying platforms [17, 18, 19] and on free-floating manipulators in simulation [20]. However, evaluation of RL on real hardware that simultaneously features thruster, reaction-wheel, *and* manipulator authority (setting in which CoM coupling is most pronounced) remains absent. Sim-to-real transfer for free-flyers with reconfigurable inertia is therefore an open problem. Our actuators are designed to be a common substrate for both communities: every actuator is fully simulated in our training environment, every characterization curve is openly released, and the same task definitions and reward structures are exposed to classical and learned controllers alike. This enables direct comparisons of, for example, a tuned LQR baseline against a PPO policy with domain randomization, on the same physical robot performing the same task; a comparison that, to our knowledge, no existing planar testbed supports end-to-end on the full thruster+RW+arms suite.

d) Contributions:

- **PINGU: platform with a set of openly released actuator modules** (a reaction wheel and two F/T-sensed robotic arms with interchangeable end-effectors) designed to co-integrate with the open-source ATMOS air-bearing testbed. We present the full mechanical and electronic design and a detailed characterization of every module (Section III), and position the resulting capabilities against the state of the art in Table I.
- **A unified ROS 2 control and estimation framework** (Section III-C) with an actuator-abstraction layer that exposes thrusters, reaction wheel, and arms through a common interface, allowing model-based (PID, LQR, MPC) and learned (PPO) controllers to be swapped without changes to the low-level stack.
- **Sim-to-real reinforcement-learning policies for dynamic disturbance rejection** under arm-induced center-of-mass and inertia shifts. We train PPO policies in IsaacLab [21] with domain randomization over base mass, center of mass, and a constant external wrench (a per-episode bias force and torque), and demonstrate transfer to hardware (Section IV).
- **Experimental evaluation across four representative tasks:** point-to-pose navigation, dynamic disturbance rejection, momentum stabilization, and force-controlled docking; comparing classical and learned controllers on the same hardware (Section IV).
- **Open-source release** of the full hardware design (CAD, BOM), firmware, ROS 2 packages, simulation environments, training scripts, and experimental logs, to make PINGU reproducible and to lower the barrier of entry to learning-based spacecraft GNC research.

The rest of the paper is organized as follows. [Section II](#) positions PINGU against prior planar emulators, manipulator-equipped testbeds, and learning-based control on real robotic platforms. [Section III](#) describes the platform design and characterization, together with the unified ROS 2 control and estimation framework. [Section IV](#) reports the experimental evaluation across the four tasks. [Section V](#) and [Section VI](#) discuss lessons learned and future directions.

II. RELATED WORK

PINGU sits at the intersection of three threads in the literature: ground-based emulators of microgravity dynamics, on-board manipulation for free-floating spacecraft, and learning-based control for real robotic systems. We survey each thread in turn, then position PINGU against the closest competitors via the comparison in [Table I](#).

A. Planar air-bearing emulators of microgravity

Air-bearing-supported floating platforms have served as mainstream terrestrial microgravity emulation for nearly half a century, and two surveys provide comprehensive historical context [[1](#), [2](#)]. The basic composition is invariant: a rigid body with porous-graphite or orifice-fed air bearings floats over an extremely flat surface made of granite, glass on optical bench, or epoxy resin, yielding three planar degrees of freedom (two translations, one rotation) with very low residual friction. Variations on this theme add hemispherical or spherical bearings to introduce limited rotational freedom, producing 5-DoF (e.g., ASTROS [[8](#)]) or 6-DoF (e.g., M-STAR [[7](#)]) facilities at the cost of mechanical complexity.

A representative cross-section of the modern planar testbed landscape illustrates the diversity of design choices. ATMOS [[5](#)], recently released as open-source by KTH, is built on a Pixhawk 6X and Jetson Orin NX stack, uses 8 PWM-driven cold-gas thrusters, and is intentionally compatible with the NASA Astrobee flight software. Bredenbeck et al. [[6](#)] describe a heavyweight (221 kg) modular stack at ESA’s Orbital Robotics Lab, equipped with 8 solenoid thrusters and a reaction wheel, and demonstrate optimal trajectory tracking for the overactuated system; the related REACSA platform [[23](#)] packages this thruster-plus-reaction-wheel actuation into a self-contained, controllable floating stack for orbital-robotics concept testing. The Luleå [[3](#)] and Luxembourg [[4](#)] floating platforms are, by contrast, lightweight (4.27 kg one and 5.0 kg the other) thruster-only design controlled by nonlinear MPC, PID and RL-based policies. At the institutional end of the spectrum, NASA JPL’s Small Satellite Dynamics Testbed [[22](#)] provides a rigorously characterized 8-thruster planar air-bearing platform whose propulsion-characterization methodology directly parallels our own ([Section III](#)). Other notable systems include the DLR TEAMS testbed for multi-spacecraft formation and contact-dynamics emulation [[24](#)], the Yonsei hardware-in-the-loop facility [[25](#)], the NUA SOOST testbed for on-orbit-operation simulation with vision-based navigation [[26](#)], the ESA ORBIT facility [[27](#)] (a 5×9 m epoxy floor with 14-camera VICON tracking), and the NTUA space-robot simulator [[28](#)]. These platforms have also become

a proving ground for increasingly sophisticated control: online active fault estimation with collision avoidance on the Caltech testbed [[29](#)], and model-predictive schemes that handle the binary, dwell-time-constrained nature of solenoid thrusters together with a reaction wheel on the DFKI/ESA free-floating platform [[30](#)] (the same actuation challenge PINGU’s base poses).

Across this body of work, two patterns can be observed. First, the prevailing actuation modality is thrusters, sometimes augmented by reaction wheels; on-board manipulators are the exception. Second, every system listed above evaluates a single controller class on a narrow task suite, making cross-platform comparison of *control strategies* difficult. PINGU retains the design characteristics established by these platforms, like the open-source Pixhawk + Jetson stack of ATMOS, and broadens it along both axes by adding a reaction wheel and two F/T-sensed manipulators within a unified actuator-abstraction layer.

B. Manipulation on free-floating ground emulators

On-orbit servicing, active debris removal, and on-orbit assembly have driven decades of work on free-flying and free-floating space manipulators, surveyed comprehensively by Flores-Abad et al. [[31](#)] and, for the dynamics and control of the free-floating regime specifically, by Moosavian and Papadopoulos [[32](#)]. The defining feature of this regime is that manipulator motion reacts back on an uncontrolled base, a coupling formalized in the foundational free-floating-manipulator literature via the generalized Jacobian’s dependence on the (shifting) mass distribution [[33](#), [34](#)], and first demonstrated in orbit by the ETS-VII free-flying space robot [[35](#)] and exploited operationally by the Orbital Express demonstration, which performed autonomous capture and robotic component transfer between two spacecraft [[36](#)]. The DLR OOS-SIM facility [[37](#)] emulates this servicer–target interaction on the ground using industrial robot arms in a hardware-in-the-loop scheme, complementing the air-bearing approach with full 6-DoF motion at the cost of true free-floating dynamics.

A smaller body of work integrates robotic manipulators directly onto air-bearing floating bases, motivated mainly by autonomous capture of free-floating targets. Sabatini et al. [[11](#)] demonstrate coordinated base-and-arm maneuvers that exploit platform motion to optimize fuel efficiency, with simulation and experimental validation. Rybus et al. [[38](#)] reproduce the operations required for an orbital capture using a free-floating manipulator mounted on a planar air-bearing simulator. Santaguida and Zhu [[9](#)] present a planar testbed with a 3-DoF manipulator, a pseudo-galactic star tracker for pose estimation, and PD trajectory tracking, validated on a tumbling-target capture scenario. Nagaoka et al. [[39](#)] show that a *dual-arm* space robot can capture and detumble a spinning target through repeated impacts, underscoring the value of two coordinated arms with contact sensing for capture — a configuration PINGU realizes with full six-axis F/T feedback at each wrist. Bob & Charlie [[10](#)] are a pair of FIT air-bearing vehicles equipped with 8-thruster RCS, a reaction wheel, and various passive grippers and grasping features — the closest existing system on actuator coverage,

TABLE I: Comparison of PINGU with existing planar/free-flying microgravity testbeds: ATMOS, Astrobee Flight (F), Astrobee Ground (G), Slider, ESA ORGL, Bob & Charlie, M-STAR (6-DoF), JPL SSDT, the UniLu ZeroG floating platform, and our platform. Entries marked – are not reported in the source paper.

	ATMOS [5]	Astrobee F [12]	Astrobee G [13]	Slider [3]	ESA [6]	Bob & Charlie [10]	M-STAR [7]	JPL SSDT [22]	ZeroG FP [4]	PINGU (Ours)
<i>Physical properties</i>										
Mass [kg]	16.80	9.58	18.97	4.27	221.67	14.90	–	–	5.32	25
Moment of inertia [kg m ²]	0.297	0.162	0.252	0.190	12.223	–	diag(1.19, 1.24, 1.43)	–	0.0591	0.72 - 1.35 ^c
Height [m]	–	0.32	0.32	–	1.025	–	–	–	0.45	0.65
Max. payload [kg]	150	–	–	–	–	–	–	–	20	1360
<i>Actuation</i>										
Num. thrusters	8	2 (impeller)	8	8	8	8	16	8	8	8
Max. thrust per axis [N]	3.4	0.849	0.849	0.7	20.0	–	–	0.51	1.0	1.5
Max. torque [N m]	1.08	0.126	0.126	–	1.5	–	–	–	–	0.24
Reaction wheel	✗	✗	✗	✗	✓	✓	✓ (×4)	✗	✗	✓
Integrated manipulators	✗ ^a	1 (perching)	1 (perching)	✗	✗	gripper only	✗	✗	✗	2 (F/T)
F/T sensing on arm	–	✗	✗	–	–	✗	–	–	–	✓
<i>Software & reproducibility</i>										
Open-source HW/SW	✓	partial	partial	✗	✗	✗	✗	✗	✗	✓
Learning-ready stack	✗	✗	✗	✗	✗	✗	✗	✗	✓ ^b	✓
<i>Demonstrated tasks</i>										
Pose tracking	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Docking / capture	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Disturbance rejection (shifting CoM)	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓
Contact-rich manipulation	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓

^a ATMOS demonstrates a manipulator *payload* but does not include it as a characterized actuator in the base release.

^b Via the external RANS/DRIFT toolchain [17], demonstrated for thruster-only control.

^c Range depending on the arms configuration.

although their manipulation hardware is passive rather than active and the platform is demonstrated only at the open-loop characterization level. The Astrobee free-flyer [12, 13], used both in flight aboard the ISS and as a ground emulator on a granite table, includes a single perching arm. A recent review of intra-vehicular free-flyers [14] surveys SPHERES, Int-Ball, CIMON, and Astrobee, concluding that despite this lineage, in-space manipulation has been demonstrated only to a limited extent, and that the dynamic interaction between a free-floating base and an active manipulator remains a central challenge for autonomy.

The collective limitation of these systems is one of *integration depth*: typically a single arm with limited sensing, no co-located reaction-wheel authority, and controllers tuned by hand for narrow tasks. None to our knowledge co-integrate dual F/T-sensed arms, a reaction wheel, and a thruster suite under a unified software stack designed to swap controllers without re-engineering the actuator interface.

C. Reinforcement learning for spacecraft and floating-platform control

Deep RL has been applied to a broad range of spacecraft GNC problems in simulation, from planetary landing [40] and asteroid proximity operations [41] to motion planning for dual-arm free-floating robots [20], with a recurring emphasis on adapting to poorly modeled dynamics (e.g., domain-randomized attitude control across two orders of magnitude of mass–inertia variation [42]). Hardware demonstrations are narrower: sim-to-real PPO for pose and velocity tracking on a thruster-only floating platform (DRIFT [17]), DRL guidance, proximity operations, and capture on the Carleton SPOT testbed [18, 43, 44], collision-free docking [19], reaction-wheel attitude control under actuator failure [45], and, most recently, a free-flyer operating aboard the ISS [46]. None of these combines thruster, reaction-wheel, and active-manipulator authority on one platform, leaving the most significant inertia-coupled case untested on real hardware.

PINGU instead draws on the broader genre of *research platforms purpose-built for learning-based control*, enabled by GPU-parallel simulation [47, 48] and validated from legged robots [49, 50] to drone racing [51] and high-speed contact tasks [52, 53]. PINGU brings this methodology to planar spacecraft emulators through an open end-to-end release, a narrow sim-to-real gap, and classical and learned controllers evaluated on the same task interfaces. It complements space-focused simulation benchmarks [16] with open hardware and focuses on reconfigurable inertia from on-board arms as the main learning challenge, which prior floating-platform RL work has not addressed.

D. Positioning of PINGU

Table I compares PINGU against representative prior platforms across actuation, sensing, software stack, and demonstrated tasks. Three observations follow. First, on *actuator coverage*, PINGU is the only system surveyed to co-integrate a thruster suite, a reaction wheel, and two F/T-sensed arms with interchangeable end-effectors; the closest prior systems cover at most two of these categories (thrusters+RW for the ESA ORGL platform and M-STAR; thrusters plus a single arm for Santaguida & Zhu [9] and Sabatini et al. [11]; thrusters+RW with passive grippers for Bob & Charlie). Second, on *software stack*, PINGU extends the open Pixhawk+Jetson+ROS 2 ecosystem shared with ATMOS by an actuator-abstraction layer spanning the heterogeneous arms+RW+thrusters mix, and is built from the outset as a substrate for both classical and learned controllers on identical task interfaces — whereas prior learning-on-hardware results in this class [17] rely on external toolchains and thruster-only actuation. Third, on *task coverage*, prior platforms typically demonstrate one or two task families; PINGU is evaluated on four (point-to-pose, disturbance rejection under shifting CoM, momentum stabilization, and force-controlled docking) with controllers swapped through a common interface, enabling a like-for-like classical-versus-learned comparison not previously reported for this platform class.




Fig. 2: **PINGU system overview.** Annotated assembly drawing of the complete platform. The chassis is organized into four functional layers from bottom to top: pneumatic base (high-pressure gas bottles and regulator manifold), actuation mid-section (eight cold-gas thrusters, reaction wheel, and LEVION mounts), power section (dual LiPo batteries and distribution electronics), and payload deck (exteroceptive sensor board and LEVION).

III. SYSTEM DESIGN AND CHARACTERIZATION

PINGU is a reconfigurable planar microgravity emulator designed for the systematic evaluation of spacecraft GNC algorithms, with particular emphasis on contact-rich proximity operations and docking build on top of the open-source platform ATMOS[5]. The complete assembled platform is shown in Fig. 2. Three co-designed subsystems are described in this section: the mechanical chassis (Section III-A), which houses the pneumatic supply, actuation modules, and payload interfaces; the electronics and sensing stack (Section III-B), which integrates flight-controller-grade avionics with a multi-modal exteroceptive sensor board; and the ROS2-based software design (Section III-C), which provides hardware abstraction and controller-deployment infrastructure. Key actuation and sensing components are characterized experimentally alongside their hardware descriptions.

A. Mechanical Architecture

PINGU is built around a modular, vertically-layered chassis measuring 40cm $\times$ 40cm $\times$ 65cm (Fig. 2). The structural skeleton consists of 20mm $\times$ 20mm aluminum T-slot extrusion profiles joined by polycarbonate side panels, while non-load-bearing enclosures, brackets, and arm linkages are fabricated via 3D printing. The chassis is organized into four stacked functional sections along its vertical axis: a pneumatic base containing the pressure bottles and regulator manifold, an actuation mid-section housing the thruster valves, reaction wheel, and robotic arm mounts, a power section with the battery packs and power-distribution electronics, and an open payload deck rated. The pneumatic supply, actuator specifications, and robotic arm configuration are detailed in the following subsections.

1) *Pneumatic Section:* Like the open-source ATMOS platform [5], PINGU uses three high-pressure gas bottles (1.5 L each, filled to 200 bar) routed through a regulator manifold on




Fig. 3: **New Way 150 mm flat round air bearing (S1015001).** Each bearing sustains a 6 aerostatic gap at 5 bar and is rated for 453.5 kg static load. Three bearings are arranged in an equilateral triangle beneath PINGU's base plate, providing near-zero friction over the test floor.

the pneumatic base plate. The manifold steps the supply down to 5 bar on two separate regulated rails: one for the three air bearings and one for the solenoid thrusters.

Frictionless planar motion is provided by three New Way [54] 150 mm flat round porous-graphite air bearings (Fig. 3), arranged in an equilateral triangle beneath the base plate. Compared to the 60 mm bearings used on ATMOS, the 150 mm format substantially increases load capacity per bearing: at the nominal 5 bar supply (within the rated 4.1 - 5.5 bar range) each bearing sustains an aerostatic gap of approximately 6 at a flow of 2.9 - 3.5 NLPM and is rated for an ideal static load of 453.5 kg. The three-bearing arrangement yields a combined bearing-limited load capacity of approximately 1360 kg, supporting a significantly heavier platform with heavier payloads, such as the LEVION and their associated hardware.

2) *Actuation Modules:* PINGU employs three complementary actuation modalities: eight cold-gas thrusters for impulsive planar force generation, a reaction wheel for continuous torque control, and two LEVION arms for inertia reconfiguration and contact manipulation. Together they cover the full range of free-floating spacecraft maneuvers within the planar microgravity environment.

a) *Thrusters:* The eight thrusters function analogously to a spacecraft's Reaction Control System (RCS): solenoid valves with a maximum switching frequency of 500 Hz, operated in software at a fixed carrier frequency of 10 Hz with a pulse width ranging from 0 ms to 100 ms (full off to full on) for proportional thrust modulation. Air is supplied at a manifold pressure of 5 bar, drawn through 4 mm tubing and expelled via 2 mm nozzles. The thrusters are arranged in a quad-corner configuration: two at each corner of the actuation mid-section with nozzle axes in mutually orthogonal body-frame directions ($\hat{x}$ and $\hat{y}$), providing independent force generation along both planar axes and torque about the yaw axis. Isolated characterization, performed by firing a single valve into a force plate, yields approximately 0.7 N per thruster. Because all eight valves share a common supply manifold, simultaneous actuation causes a pressure drop that reduces individual thrust output; the resulting coupling is characterized in Fig. 4 and

must be accounted for in any thrust-allocation model.




Fig. 4: Thruster characterization. *Left:* Measured force output as thrusters are opened sequentially (1 through 8), showing the progressive pressure drop across the shared supply manifold that reduces individual thrust output under concurrent actuation. *Right:* CAD model of the thruster box alongside the assembled unit mounted on the PINGU actuation mid-section.

b) Reaction wheel: A single reaction wheel provides continuous, non-impulsive attitude authority, complementing the binary solenoid thrusters which cannot produce smooth yaw torque profiles. The wheel is a 20 cm diameter disk 3D-printed in metal-blended PLA to maximise rotational inertia at low mass; the open-source design can equally be CNC-machined from any solid material for higher inertia requirements. It is driven by a D5312S 330 KV dual-shaft brushless motor controlled by an ODrive S1 [55], with the disk mounted on one output shaft and an AMT212B magnetic absolute encoder on the other. The motor delivers a torque constant of $0.025 \frac{Nm}{A}$. The reaction wheel angular velocity Ω_r and generated torque τ_r are modelled in discrete time as

$$\Omega_{r,k+1} = \frac{\tau_{cmd,k}}{b} + \left(\Omega_{r,k} - \frac{\tau_{cmd,k}}{b} \right) \exp \left(-\frac{b}{J_r} \Delta t \right), \text{ and}$$

$$\tau_{r,k+1} = J_r \dot{\Omega}_{r,k+1} = (\tau_{cmd,k} - b \Omega_{r,k}) \exp \left(-\frac{b}{J_r} \Delta t \right)$$

respectively, where b is a viscous friction constant, J_r is the wheel moment of inertia about its spin axis, and τ_{cmd} is the command torque at a given step. Fig. 5 show a comparison between the modelled wheel dynamics and the measured torque-speed envelope, across the full operating speed range.

c) LevionArms: The LEVION are the most distinctive element of PINGU: a pair of planar 2-DOF robotic arms that serve a dual role not found in existing air-bearing testbeds. First, by repositioning their links, they actively redistribute the platform's effective centre of mass and moment of inertia, enabling systematic evaluation of controller robustness across different inertial regimes without changing the hardware. Second, they act as contact-manipulation end-effectors for proximity operations and docking, exercising the full sensing and actuation stack in a single closed-loop task.

Each arm forms a shoulder-elbow serial chain (Fig. 6) with a total mass of 2.28 kg, with links 3D-printed for low mass and rapid reconfiguration. Both joints are actuated by CubeMars AK80-8 KV60 quasi-direct-drive actuators (rated

output torque: 10 Nm; peak output torque: 25 Nm; 8:1 reduction ratio). During the experiments, the commanded joint velocity and torque were limited to 10 rad/s and 10 Nm, respectively; the low gear ratio maximizes backdrivability and torque transparency, which is critical for compliant contact. The shoulder joints share a symmetric range of ± 1.56 rad; the elbow joints are geometrically mirrored — left elbow $[-2.61, -0.36]$ rad, right elbow $[0.36, 2.61]$ rad — yielding a manipulation workspace symmetric about the platform centreline. Each arm's motors are torque-controlled through a PD law that allows compliant position reference tracking. A Leptrino six-axis F/T sensor at each wrist provides full wrench feedback for impedance control and contact-onset detection. A standardized wrist coupling enables rapid interchange between grippers, suction cups, and custom tools without any structural modification.

3) Payload: PINGU rests on three New Way 150 mm flat round air bearings, each rated at 453.5 kg, giving a combined bearing-limited load capacity of 1360 kg; in practice, the usable payload budget is constrained by the platform's own mass and the laboratory air supply. The mid-deck section provides side mounts for the two LEVION arms, leaving the top deck as a dedicated, highly customizable payload area. In the default configuration, this top deck carries the exteroceptive sensor board. The entire mounting interface is designed to easily accommodate alternative actuation modules, sensor arrays, or specialized experimental hardware, allowing the platform to be reconfigured for different tasks and weight-distribution requirements without structural modifications.

B. Electronics and Sensors

PINGU's electronics subsystem integrates a flight-grade compute stack, a shared CAN actuator bus, and a dual-battery power distribution network, together with a dedicated exteroceptive sensor payload described in Section III-B2.

1) Electronics: The central compute stack is the Holybro Pixhawk Jetson Baseboard, which integrates a Pixhawk 6X flight controller and an NVIDIA Jetson Orin NX (16 GB RAM) into a single package, simplifying the hardware and software integration of PX4 with a companion computer. The Pixhawk and Jetson communicate via the board's internal




Fig. 5: Reaction wheel. (A) Measured torque-speed envelope of the reaction wheel assembly, showing usable attitude-control authority across the operating speed range. (B) The 20 cm metal-blended PLA disk mounted on the D5312S dual-shaft motor inside the PINGU actuation mid-section.




Fig. 6: **LEVION design.** Planar 2-DOF arm mounted on PINGU’s actuation mid-section. Each joint is driven by a CubeMars AK80-8 KV60 quasi-direct-drive motor; a Leprino six-axis F/T sensor at the wrist provides wrench feedback for impedance control. The symmetric bilateral mounting allows the arms to reconfigure the platform’s effective inertia distribution on-the-fly.

Gigabit Ethernet switch, which carries both the ROS 2 inter-process traffic and the micro-XRCE-DDS bridge between PX4 and the companion computer.

a) *Actuator drivers and CAN bus:* The reaction wheel motor and all four joints of the two LEVION arms share a single CAN bus routed through the actuation mid-section. The reaction wheel is driven by an ODrive S1 motor driver; each arm joint is driven by the integrated CAN-FD interface of the CubeMars AK80-8 KV60 actuator. All five joints are exposed to ROS 2 via a `ros2_control` hardware interface running at 200 Hz, providing a hardware-agnostic actuation API to the software stack.

b) *Power distribution:* The system is powered by two 6S 9.5 Ah LiPo batteries (22.2 V nominal) with a shared ground, partitioned into two isolated domains to prevent high-current actuation transients from disturbing the control electronics (Fig. 7). The *control domain* draws from the first battery through a regulated 24 V buck-boost converter, supplying the Pixhawk 6C, Jetson Orin NX, sensor board, and thruster solenoids. The *actuation domain* draws from the second battery directly, powering the ODrive S1 motor controller and all four CAN-FD arm actuators.

2) *Sensor Suite:* PINGU’s sensing is organized across three complementary layers. Proprioceptive state is provided by the Pixhawk 6C flight controller, which exposes a 3-axis IMU (accelerometer and gyroscope), barometer, and magnetometer as standard ROS 2 sensor topics via the micro-XRCE-DDS bridge, making inertial measurements available to all control nodes without additional drivers. Each LEVION arm wrist embeds a 6-axis force/torque sensor whose readings are surfaced through the `ros2_control` hardware interface, providing direct contact sensing for manipulation and docking tasks. Exteroceptive perception is offloaded to a dedicated sensor board based on the NVIDIA Jetson AGX Orin Developer Kit, mounted on PINGU’s payload deck; all sensor drivers run on this board and publish data as ROS 2 topics available to the main Jetson Orin NX over the onboard Ethernet network (Fig. 8).

The exteroceptive suite comprises four complementary modalities. An Intel RealSense D455 RGB-D camera provides colour imagery and depth point clouds via an active infrared stereo pair, primarily used for close-range target detection and




Fig. 7: **Electronics overview.** Two-domain power architecture of PINGU. The control domain (Battery 1, 22.2 V nominal, regulated) supplies the flight controller, compute, and solenoid loads. The actuation domain (Battery 2, direct) supplies the ODrive S1 and CAN-FD arm actuators. Signal paths – including CAN bus, micro-XRCE-DDS bridge, and Ethernet – are shown alongside the power rails. Both domains share a common ground.




Fig. 8: **PINGU sensor suite.** The four exteroceptive modalities mounted on the dedicated sensor board (Intel RealSense D455, FLIR Firefly, Prophesee event camera, Livox LiDAR), together with representative output from each.

docking. A FLIR Firefly RGB monocular camera provides a high-frame-rate colour stream. A Prophesee event camera offers sub-millisecond-latency asynchronous brightness-change detection, enabling high-speed proximity sensing robust to motion blur. A Livox mid-range LiDAR provides dense 3-D point clouds for environment mapping and obstacle avoidance. Together, the four modalities span the full spectrum from high-frequency, low-latency event data to dense geometry, allowing sensing pipelines to be benchmarked on the same platform and task suite.

C. Software Design

PINGU’s software is organized into two cooperating ROS 2 workspaces that cleanly separate hardware abstraction from control logic (Fig. 9). The low-level platform workspace runs on the Jetson Orin NX and is responsible for exposing every actuator as a standard ROS 2 interface. The controller-deployment workspace runs on the same machine and implements all control and inference logic. This separation allows controllers (classical or learned) to be developed, evaluated,




Fig. 9: Two-workspace ROS2 software architecture of PINGU: the low-level platform workspace exposes all actuators as standard interfaces; the controller-deployment workspace implements classical and learned control logic above that abstraction boundary.

and swapped without any modification to the hardware-facing layer.

a) *Low-level hardware abstraction*: All actuators are exposed through two complementary mechanisms. Thruster commands are handled by a dedicated PX4 offboard node that places the Pixhawk 6X into direct-actuator offboard mode via the micro-XRCE-DDS bridge and forwards an 8-element command vector as an `ActuatorMotors` uORB message over the DDS bus. The robotic arms and reaction wheel are managed by a `ros2_control` hardware interface that communicates with the CAN bus at 200 Hz, exposing position, velocity, and effort command interfaces for all five actuated joints. A command multiplexer node arbitrates between concurrent command sources, allowing a user to manually control the platform through a joystick or let the autonomous controllers seamlessly take over, while enforcing priority ordering and safety limits across the two actuation paths. Sensor data enters the ROS 2 graph through two drivers: a VRPN client publishes OptiTrack motion-capture poses at 100 Hz, and a RealSense driver streams RGB-D frames for vision-based tasks.

b) *Controller deployment and inference*: All controllers share a common output interface: a 15-element command vector comprising a bearing flag, a thruster-enable bit, eight binary thruster duty-cycle values, four LEVION arm joint setpoints, and one reaction-wheel torque command. Classical controllers are implemented as ROS 2 nodes that subscribe to the state topics and publish to this interface directly. Learned policies are deployed via a ROS 2 node built on the RoboRAN framework [15], which runs a fixed-rate inference loop at 10 Hz (Fig. 10). The node composes four interchangeable modules in a fixed pipeline. A *state preprocessor* converts raw sensor data into the body-frame state required by the policy; the default implementation buffers OptiTrack pose messages at 100 Hz over a 30-frame window and numerically differentiates them to obtain linear and angular velocities, while drop-in alternatives (ArUco marker detection, odometry) are registered via the same factory interface. A task-specific *observation formatter* assembles the policy input tensor from




Fig. 10: **ROS2 deployment node graph**. External inputs (red) feed the *Task Node*, which chains four interchangeable modules (blue): a *state preprocessor* (sensor drivers and velocity estimation), *observation formatter* (task-specific input assembly), *inference runner* (policy evaluation), and *robot interface* (action mapping). The robot interface publishes a unified command to the *CMD Mux*, which arbitrates between autonomous and manual (joystick) control sources. The *Low-Level Controller* (green) routes thruster commands to PX4 via an `ActuatorMotors` uORB message and joint commands to `ros2_control` (yellow), which manages effort, position, velocity, and trajectory controllers for the reaction wheel and both LEVION arms over the CAN bus. Every module is hot-swappable at launch time; the hardware layer requires no modification when the task, sensor, or policy changes.

the preprocessed state and a received goal; for example, the *GoToPose formatter* computes goal distance, heading errors (sine/cosine encoded), body-frame velocities, and the previous action vector. An *inference runner* evaluates the loaded policy at each step. The platform natively supports RSL-RL [56] checkpoints, executing both MLP and GRU architectures while directly restoring empirical observation-normalization statistics. For broader compatibility, the runner also integrates ONNX Runtime (leveraging GPU acceleration via NVIDIA TensorRT on the Jetson), SKRL, RL-Games, and a classical LQR baseline. A *robot interface* maps the policy output to the shared 15-element command vector. A simulation bridge is also provided, allowing a policy to step inside IsaacLab [21] while issuing live actions to the hardware, enabling closed-loop validation before committing to fully off-policy deployment. Switching between a classical and a learned controller requires only relaunching the controller-layer nodes; the hardware-abstraction layer and all sensor pipelines remain unchanged.

c) *Digital Twin*: The PINGU digital twin is implemented within the Space Robotics Bench (SRB) [16], a GPU-accelerated robot-learning framework built on top of IsaacLab [57] and NVIDIA Isaac Sim. The simulated environment faithfully reproduces the full platform dynamics — rigid-body kinematics, binary solenoid-valve thrusters, reaction wheel, and 2-DOF LEVION arms with configurable end-effectors — as well as the physical laboratory layout, including the precision-levelled air-bearing test floor, an OptiTrack motion-capture system, and two ceiling-mounted Universal Robots UR10 manipulators that serve as interactive targets



Fig. 11: **PINGU digital twin in the Space Robotics Bench.** Thousands of parallelized copies of the laboratory environment running simultaneously on a single GPU inside IsaacLab. Each instance contains a full PINGU model (air-bearing platform with two LEVION arms) and two UR10 manipulator mounted to the ceiling and wall rails, replicating the physical laboratory layout. Massively parallel rollouts with domain randomization enable the sim-to-real transfer results reported in Section IV.

for proximity-operations and contact-dynamics tasks (Fig. 11). Exploiting PhysX rigid-body simulation, SRB runs thousands of environment instances in parallel on a single GPU, enabling sample-efficient RL training with domain randomization. Beyond training, the digital twin is used for controller validation before hardware deployment and as a real-time state-visualization interface during laboratory experiments.

IV. EXPERIMENTAL EVALUATION

The experiments presented here are designed to demonstrate the breadth of PINGU’s capabilities rather than to benchmark algorithms in isolation. Concretely, we aim to show three things: (i) the platform’s unified actuator-abstraction layer allows both classical optimal controllers and end-to-end learned policies to be deployed on the same hardware, on the same tasks, without modifying the underlying stack; (ii) the platform’s modular sensor suite can be integrated directly into the control pipeline, enabling sensor-driven tasks beyond blind pose tracking; and (iii) the digital twin is a faithful enough model of the real platform that policies trained entirely in simulation transfer to hardware with acceptable performance.

a) *Task overview:* The four tasks are organized in increasing order of complexity. **Point-to-Pose Navigation** is the foundational planar maneuver, evaluated with both a classical LQR and a learned PPO controller across three actuator subsets (thrusters only, thrusters + reaction wheel, and all actuators), establishing baseline accuracy and a direct classical-vs-learned comparison across the three arm configurations of Fig. 12. **Dynamic Disturbance Rejection** takes the same navigation task and adds a continuous, unobservable CoM disturbance driven by the LEVION arms oscillating autonomously, isolating the learned policy’s ability to reject structured, low-frequency inertia shifts that a rigid-body model cannot fully capture. **Momentum Damping / Stabilization** isolates the reaction wheel alone and evaluates a PPO policy’s ability to damp arbitrary initial angular velocities across the three




Fig. 12: **LEVION arm configurations.** Three static arm poses used throughout the experiments to vary the platform’s effective moment of inertia: *Rest* (arms extended outward at 90° , maximizing inertia); *Closed* (arms tucked toward the body, minimizing inertia); and *Side* (both arms extended to one side, producing an asymmetric mass distribution). The configurations are held fixed during each run and are not part of the controller’s action space; they serve as a structured way to stress-test controller robustness across different inertial regimes.

arm configurations, probing the wheel’s real authority under varying effective inertia. **Force-Controlled Docking** is the most integration-intensive task: an ArUco marker detected by the onboard RGB-D camera provides a live 6-DoF goal to the learned point-to-pose controller, and contact with a fixed wall is managed through joint-space impedance control on the LEVION arms, with contact onset detected by the wrist F/T sensors — exercising the full sensing and actuation stack in a single closed-loop pipeline.

A. Tasks

1) *Point-to-Pose Navigation:* In this task, PINGU must navigate from a random initial pose to a fixed target pose (x^*, y^*, ψ^*) . This is the platform’s foundational maneuver and the first task on which both classical and learned controllers are evaluated side by side: every other task in this evaluation either builds directly on it (point-to-pose under disturbance, vision-guided docking) or reuses its reward structure. To highlight the system’s reconfigurability, the task is executed using three actuator subsets: *thrusters only*, *thrusters and reaction wheel*, and *all actuators* (including arms). For the two reduced subsets, the LEVION arms are held static at one of the three fixed poses introduced in Fig. 12 (*Side*, *Rest*, *Closed*), so that the effective inertia is varied independently of the controller. A classical LQR controller (with integral action over position and heading) and a learned PPO controller are both deployed on hardware for the *thrusters only* and *thrusters + reaction wheel* subsets; PPO alone is evaluated on the *all-actuators* configuration since the arm joints are part of its action space. Performance is quantified by the final position error e_p [m] and heading error e_o [rad]. Results across all configurations




Fig. 13: **Point-to-pose navigation.** Left: hardware trajectories of PINGU navigating from four random initial poses (colored dots) to a fixed target pose and heading (arrow), with arrowheads marking direction of travel. Right: corresponding distance error (top) and heading error (bottom) over time for the same four runs, both converging to near zero well within the 60 s trial.

are reported in Table II.

Because this reward is reused, unmodified, by the disturbance-rejection task and underlies the navigation component of the docking task, we define it in full here. For the learning controllers, the reward function $r_t^{\text{pose-nav}}$ is

$$r_t^{\text{pose-nav}} = r_t^{\text{pose}} + r_t^{v,\omega} + r_t^{\text{bnd}}$$

where r_t^{pose} rewards the robot for reaching the target position and heading simultaneously, $r_t^{v,\omega}$ penalizes the robot for exceeding a desired range of linear and angular velocities, and r_t^{bnd} penalizes the robot’s proximity to a defined maximum distance boundary, keeping it within the workspace. The full reward expression and weights are given in Table VII.

We demonstrate the PINGU platform’s ability to host both model-based and learned controllers through the same actuator-abstraction layer without any hardware modification by successfully deploying both PPO [58] and LQR[59], and showing that converge to the target pose. The PPO controller consistently outperforms LQR on position accuracy across all actuator subsets and arm configurations: in the *thrusters-only* configuration, PPO reduces e_p by 55–69% relative to LQR (e.g., 0.024 m vs. 0.083 m for Side, 0.053 m vs. 0.170 m for Closed); in the *thrusters + reaction wheel* configuration, PPO similarly halves the position error (Rest: 0.048 m vs. 0.109 m; Closed: 0.068 m vs. 0.130 m). Heading errors are more comparable between the two controllers, with LQR occasionally matching PPO on that metric alone. The *all-actuators* configuration, available to PPO only, achieves the lowest real-world position error overall ($e_p = 0.0081$ m), reflecting the additional control authority the arms provide when included in the action space. The sim-to-real gap for

<i>Point-to-Pose Navigation</i>		
Methods	$e_p[\text{m}] \downarrow$	$e_o[\text{rad}] \downarrow$
(a) All actuators		
Sim	0.0022 (± 0.002)	0.0044 (± 0.003)
Sim + DR	0.0124 (± 0.019)	0.0129 (± 0.012)
PPO (Real)	0.0081 (± 0.001)	0.0257 (± 0.015)
(b) Thrusters and reaction wheel		
<i>Rest</i>		
Sim	0.0220 (± 0.022)	0.0047 (± 0.004)
Sim + DR	0.0097 (± 0.010)	0.0069 (± 0.006)
PPO (Real)	0.0482 (± 0.027)	0.0279 (± 0.020)
LQR (Real)	0.1088 (± 0.084)	0.0630 (± 0.095)
<i>Closed</i>		
Sim	0.0045 (± 0.006)	0.0040 (± 0.003)
Sim + DR	0.0145 (± 0.011)	0.0067 (± 0.007)
PPO (Real)	0.0679 (± 0.050)	0.0598 (± 0.035)
LQR (Real)	0.1303 (± 0.076)	0.1005 (± 0.118)
<i>Side</i>		
Sim	0.0137 (± 0.011)	0.0047 (± 0.004)
Sim + DR	0.0083 (± 0.009)	0.0081 (± 0.009)
PPO (Real)	0.0547 (± 0.042)	0.0275 (± 0.015)
LQR (Real)	0.097 (± 0.057)	0.0306 (± 0.011)
(c) Thrusters only		
<i>Rest</i>		
Sim	0.0057 (± 0.009)	0.0046 (± 0.003)
Sim + DR	0.0056 (± 0.007)	0.0105 (± 0.008)
PPO (Real)	0.0745 (± 0.040)	0.0219 (± 0.003)
LQR (Real)	0.0843 (± 0.036)	0.0461 (± 0.067)
<i>Closed</i>		
Sim	0.0099 (± 0.011)	0.0055 (± 0.004)
Sim + DR	0.0075 (± 0.009)	0.0062 (± 0.005)
PPO (Real)	0.0527 (± 0.019)	0.0414 (± 0.036)
LQR (Real)	0.1696 (± 0.084)	0.0307 (± 0.016)
<i>Side</i>		
Sim	0.0064 (± 0.011)	0.0024 (± 0.003)
Sim + DR	0.0069 (± 0.009)	0.0087 (± 0.008)
PPO (Real)	0.0241 (± 0.016)	0.0335 (± 0.016)
LQR (Real)	0.0830 (± 0.025)	0.0229 (± 0.023)

TABLE II: **Point-to-pose navigation.** Final position error e_p and heading error e_o comparing the learned PPO controller (Sim, Sim+DR, and Real) against the classical LQR controller (Real only), across three actuator subsets and three static LEVION arm poses (Side, Rest, Closed). The all-actuators configuration has no LQR baseline as the arms are part of the PPO action space.

PPO is consistent across the three static arm poses within a given actuator subset, suggesting it is driven by unmodeled thruster and air-bearing dynamics rather than by the specific arm configuration.




Fig. 14: **Dynamic disturbance rejection.** Top: heading error θ over time for four representative hardware runs, converging toward zero despite the continuous CoM disturbance. Middle: corresponding distance error to the target pose, converging with bounded residual oscillations driven by the disturbance. Bottom: commanded shoulder/elbow joint positions for the same four runs, showing the disturbance generator’s continuous ± 0.8 rad oscillation throughout each 60 s trial. Frame sequence: a representative run, with the platform navigating to the target pose while its arms swing autonomously.

2) *Dynamic Disturbance Rejection (CoM Shifting)*: In this task, PINGU must navigate to a target pose (x^*, y^*, ψ^*) while its two LEVION arms introduce continuous, uncontrolled disturbances to the system’s CoM. The arms oscillate autonomously between ± 0.8 rad at the shoulder joints through a four-state finite state machine (hold right $\rightarrow$ swing left $\rightarrow$ hold left $\rightarrow$ swing right), with each transition duration sampled uniformly from $[0.1, 0.3]$ s and cosine-interpolated to produce smooth, time-varying CoM displacements. Because the arms are driven by a separate position controller and are not part of the policy’s observation space, their motion constitutes an exogenous disturbance that the agent must reject implicitly through its thruster and reaction-wheel commands alone. The task otherwise follows the same reward function as the Point-to-Pose Navigation task. It specifically evaluates the agent’s ability to reject a persistent, structured CoM perturbation without direct observability of its source, probing the robustness of the learned momentum-management strategy.

The heading error degrades only mildly from simulation to hardware: domain randomization alone raises e_o from 0.0074 rad (Sim) to 0.0340 rad (Sim + DR), and the real-hardware result ($e_o = 0.0346$ rad) is nearly indistinguishable from the Sim + DR value, indicating that domain randomization closes most of the heading sim-to-real gap. Position tracking is harder to transfer: the real position error ($e_p = 0.0889$ m) is more than five times larger than the

Methods	<i>Dynamic Disturbance Rejection</i>	
	$e_p[m] \downarrow$	$e_o[rad] \downarrow$
Sim	0.0070 (± 0.005)	0.0074 (± 0.011)
Sim + DR	0.0156 (± 0.016)	0.0340 (± 0.057)
Real	0.0889 (± 0.041)	0.0346 (± 0.018)

TABLE III: **Dynamic disturbance rejection.** Final position error e_p and heading error e_o for the PPO policy navigating to a fixed target pose under continuous, unobserved CoM disturbance from arm motion, evaluated in simulation, in simulation with domain randomization (DR), and on hardware.

Sim + DR value ($e_p = 0.0156$ m), a substantially larger jump than for heading. Because the policy never observes the arms directly, it must infer and reject the resulting CoM-induced positional drift purely from the platform’s own pose feedback, so residual modeling mismatches in the real thruster and air-bearing dynamics are amplified by the disturbance’s persistent, never-settling nature. Despite this gap, the policy maintains stable closed-loop tracking throughout each 60 s trial, with the bounded distance-error oscillations visible in Fig. 14 confirming that it successfully rejects a persistent, unobserved CoM perturbation without diverging.

3) *Momentum Damping / Stabilization*: In this task, PINGU is initialized with a random angular velocity drawn uniformly from $[\pm 0.1, \pm 0.5]$ rad/s and must bring its rotation to rest using the reaction wheel alone, with no thruster actuation. The experiment probes the reaction wheel’s angular-momentum absorption capacity across three LEVION arm configurations that vary the platform’s effective moment of inertia (Fig. 12): *Side* (producing an asymmetric mass distribution); *Rest* (extended radially outward and maximizing effective inertia); and *Closed* (minimizing effective inertia). A PPO policy is trained to command the reaction wheel torque as a function of the current angular velocity; Fig. 15 illustrates the resulting angular-velocity decay and commanded torque for four representative runs in the *Rest* configuration. Performance across all three configurations is quantified by the residual angular velocity ϕ [rad/s] and the half-time $t_{1/2}$ [s], the time for the angular velocity to reach half its initial magnitude, as reported in Table IV.

In simulation, the policy stabilizes the platform at a comparable rate and residual velocity across all three arm configurations ($t_{1/2} \approx 3.8$ – 4.1 s, $\phi \approx 0.043$ – 0.047 rad/s), indicating low sensitivity to the arm-induced inertia changes considered here. On hardware, the policy decays the angular velocity fastest in the *Side* configuration ($t_{1/2} = 4.845$ s), as also illustrated by the representative *Rest*-configuration runs in Fig. 15, but settles to the largest residual velocity of the three ($\phi = 0.0594$ rad/s); conversely, *Rest* decays more slowly ($t_{1/2} = 6.377$ s) yet reaches the lowest residual velocity overall ($\phi = 0.0228$ rad/s), while *Closed* performs worst on both metrics ($t_{1/2} = 6.580$ s, $\phi = 0.0534$ rad/s). This rank reversal between fastest initial decay and lowest long-term residual velocity, absent from the simulated results, indicates that the reaction wheel’s real disturbance-rejection authority interacts with the platform’s




Fig. 15: **Momentum damping in the Rest arm configuration.** Top: magnitude of the platform’s angular velocity over time across four representative hardware stabilization runs; dotted lines mark half the initial angular velocity and diamond markers indicate each run’s half-time $t_{1/2}$. Bottom: corresponding commanded reaction-wheel torque for the same four runs, showing the PPO policy applying progressively less braking torque as the angular velocity decays toward zero.

Methods	Stabilization	
	$\phi[\text{rad/s}] \downarrow$	$t_{\text{half}}[\text{s}] \downarrow$
(b) Rest		
Sim	0.0467 (± 0.015)	3.7975 (± 1.607)
Real	0.0228 (± 0.023)	6.377 (± 4.941)
(c) Closed		
Sim	0.0432 (± 0.015)	3.8773 (± 1.646)
Real	0.0534 (± 0.093)	6.580 (± 5.046)
(a) Side		
Sim	0.0429 (± 0.015)	4.0773 (± 4.750)
Real	0.0594 (± 0.030)	4.845 (± 1.873)

TABLE IV: **Momentum damping / stabilization.** Residual angular velocity ϕ and half-time $t_{1/2}$ of the reaction-wheel-only PPO stabilization policy, evaluated in simulation and on hardware across three LEVION arm configurations (Side, Rest, Closed) spanning the platform’s range of effective moment of inertia.

effective inertia distribution in ways not fully captured by the simulated dynamics, particularly in the low-velocity tail of the response.

4) *Force-Controlled Docking:* This task showcases the integration of PINGU’s full sensing and control stack into a single pipeline, chaining vision-based pose estimation, learned navigation control, and contact-triggered compliant control across three of the platform’s actuation modalities. The Intel RealSense D455 RGB camera on the sensor payload detects an ArUco fiducial marker affixed to a fixed wall and estimates its 6-DoF pose relative to the platform in real time; this pose




Fig. 16: **Force-controlled docking.** Top: frame sequence of a representative run, with the platform’s trajectory approaching the wall-mounted ArUco target. Middle: magnitude of the F/T-sensed contact force at the left (solid) and right (dashed) wrist over four repeated runs, showing intermittent contact spikes once the arms reach the wall. Bottom: magnitude of the platform’s linear velocity for the same four runs, peaking mid-approach and decaying to near zero as the arms’ compliant contact dissipates the platform’s residual kinetic energy.

is fed as the goal to the same learned point-to-pose controller evaluated above, which drives the floating platform toward the wall. As the LEVION arms make contact with the wall, the Leprino six-axis F/T sensor at each wrist registers the contact event. Rather than actively braking with the thrusters or reaction wheel, PINGU lets its own inertia carry the residual approach velocity into the wall, and the resulting contact is damped entirely at the arm level by a joint-space impedance controller running in the LEVION arms’ `ros2_control` loop: each joint tracks a fixed position setpoint through a virtual spring-damper law with fixed K_p and K_d gains, so that the arm yields compliantly on contact and dissipates the platform’s residual kinetic energy without letting the wall’s reaction force push the floating base back away from the docking surface. Fig. 16 reports results over multiple docking runs, showing the F/T force measured at the left and right wrists together with the platform’s velocity at the end of each contact.

V. DISCUSSION

The experiments of Section IV were designed less to crown a single best controller than to demonstrate that PINGU can host a heterogeneous set of controllers and tasks on one chassis. Several lessons emerge from running them.

a) *Classical and learned control share the same hardware, with complementary strengths:* The headline systems result is that a hand-tuned LQR and an end-to-end PPO policy both deploy on PINGU through the identical actuator-abstraction layer, on identical tasks, with no change to the low-level stack (Section IV-A1). Within that common envelope the two controllers exhibit a clear and consistent trade-off: the learned policy roughly halves the final position error of

LQR across every actuator subset and arm configuration, while the two are far closer on heading, where LQR occasionally matches PPO. The advantage widens when the arms are admitted into the policy’s action space — the all-actuators configuration attains the lowest position error in the study ($e_p = 0.0081\text{ m}$), confirming that the additional control authority of the arms is usable by a learned controller in a way an LQR built on a fixed rigid-body model cannot exploit. This is precisely the kind of like-for-like comparison PINGU is built to enable, and it is, to our knowledge, the first reported across the full thruster + reaction-wheel + arm envelope on real hardware.

b) Real attitude authority depends on effective inertia in ways simulation under-predicts: The reaction-wheel stabilization task (Table IV) is the clearest example of a phenomenon that only the hardware reveals. In simulation the policy damps the platform at near-identical rates across all three arm configurations; on hardware a rank reversal appears between fastest initial decay (*Side*, $t_{1/2} = 4.85\text{ s}$) and lowest long-term residual velocity (*Rest*, $\phi = 0.0228\text{ rad/s}$), with the minimum-inertia *Closed* pose worst on both metrics. The reaction wheel’s true disturbance-rejection authority therefore interacts with the platform’s mass distribution and with friction in the low-velocity tail in a manner the simulated dynamics do not fully capture. That this surfaces only on hardware underscores the value of a platform on which effective inertia can be varied as a controlled, repeatable experimental variable rather than a modeling assumption.

c) Where the sim-to-real gap lives: Across tasks the gap is strongly anisotropic. Domain randomization closes almost all of the *heading* gap — in disturbance rejection the real heading error (0.0346 rad) is nearly indistinguishable from the Sim + DR value (0.0340 rad) — whereas *position* tracking transfers far less cleanly, with the real error exceeding Sim + DR by more than fivefold under a persistent CoM disturbance. Two observations localize the cause. First, the gap is consistent across the three static arm poses within a given actuator subset, which points away from the arm-induced inertia change and toward unmodeled thruster and air-bearing dynamics as the dominant residual. The shared-manifold pressure coupling characterized in Fig. 4, whereby simultaneous valve firing depresses per-thruster force, is a prime suspect: it directly corrupts the translational impulse budget while leaving differential yaw commands comparatively intact. Second, the low-rate 10 Hz policy-inference loop, set against 100 Hz state feedback and a 200 Hz joint loop, bounds how quickly translational drift can be corrected once it accumulates. Both are concrete, characterizable targets — manifold-aware thrust allocation and a higher-rate inference path — rather than fundamental limits.

d) Reconfigurable inertia as a first-class research instrument: The LEVION arms serve two distinct roles in these experiments, and both proved essential. Held static in the three poses of Fig. 12, they let us sweep effective inertia without touching the hardware; driven autonomously, they generate a continuous, unobserved CoM disturbance that the navigation policy must reject blind (Fig. 14). The policy maintains bounded, non-divergent tracking throughout each trial despite

never observing the arms — evidence that learned controllers can absorb the structured, hard-to-model inertia variation that classical model-based design handles least gracefully. This is the central capability distinguishing PINGU from thruster-and-wheel-only emulators, and it is what makes the platform a substrate for the open problem of sim-to-real RL under reconfigurable inertia.

e) Integration and maintainability: Finally, the experiments are themselves evidence for the system-level design choices. The two-workspace ROS 2 architecture (Fig. 9) meant that swapping LQR for PPO, or a fixed marker goal for a live RGB-D detection in the docking task, required relaunching only controller-layer nodes; the hardware-abstraction and sensor layers were untouched across all four tasks. The two-domain power split kept high-current arm and wheel transients off the control rails over long experiment campaigns. These are not headline scientific results, but they determine whether a platform survives contact with real research use — and they are the properties we most hope the open-source release transfers to other groups.

VI. CONCLUSION

We presented PINGU, a reconfigurable planar microgravity platform that co-integrates eight cold-gas thrusters, a reaction wheel, and two force/torque-sensed LEVION arms with interchangeable end-effectors that extends the open-source ATMOS testbed, under a unified ROS 2 stack that exposes every actuator through a common abstraction layer. We characterized each subsystem, described the control and simulation infrastructure, and validated the platform on four tasks — point-to-pose navigation, dynamic disturbance rejection under arm-induced CoM shifts, reaction-wheel stabilization, and force-controlled docking — deploying classical and learned controllers on the same hardware through the same interface. Taken together, the results support the paper’s central claim: PINGU is less a single result than an *enabler*, a common substrate on which the spacecraft-GNC and reinforcement-learning communities can pose and answer questions that no existing planar testbed supports end-to-end.

That framing is clearest when set against the representative use cases of Fig. 1. Of the five envisioned, three — docking, point navigation, and momentum stabilization — are demonstrated here; the remaining two define our immediate roadmap.

a) Dual-arm capture and detumbling: The *target-catching* use case is the natural next step for the arms: coordinated dual-arm capture and detumbling of a free, tumbling target, with the wrist F/T sensors closing the compliant-impact loop that contact-rich capture demands [39, 44, 9]. The ceiling- and wall-mounted UR10 manipulators already modeled in the digital twin (Fig. 11) provide a repeatable, programmable moving target for exactly this scenario, letting capture policies be trained in simulation and transferred to a controlled physical rehearsal.

b) Magnetic wall locomotion: The *wall-locomotion* use case exploits PINGU’s standardized wrist coupling: fitting magnetic end-effectors lets the arms “walk” the platform along a ferromagnetic structure, a propellant-free locomotion and

station-keeping mode for inspection and traversal of large structures. This recasts the arms from manipulators into a locomotion subsystem and stresses the base–arm coupling in a regime — sustained contact under intermittent anchoring — that none of our current tasks exercises.

c) Richer perception in the loop: The docking experiment relied on a single ArUco marker and the RGB-D camera; the platform’s event camera, LiDAR, and global-shutter RGB sensor remain outside the control loop. Bringing them in — for markerless, high-rate relative pose estimation of non-cooperative targets — is the perception counterpart to the capture and locomotion goals above, and the multi-modal sensor board was provisioned with exactly this progression in mind.

We release the full hardware design, firmware, ROS 2 packages, simulation environments, and experimental logs to lower the barrier to entry for learning-based spacecraft GNC research, in the hope that PINGU serves the community as a shared, extensible foundation for the contact-rich, inertia-coupled proximity-operations problems that remain open.

ACKNOWLEDGMENTS

The authors would like to thank the ATMOS team for their invaluable help and support throughout this work.

REFERENCES

- [1] Jana L. Schwartz, Mason A. Peck, and Christopher D. Hall. Historical review of air-bearing spacecraft simulators. *Journal of Guidance, Control, and Dynamics*, 26(4):513–522, 2003. doi: 10.2514/2.5085.
- [2] Tomasz Rybus and Karol Seweryn. Planar air-bearing microgravity simulators: Review of applications, existing solutions and design parameters. *Acta Astronautica*, 120: 239–259, 2016.
- [3] Avijit Banerjee, Sumeet G. Satpute, Christoforos Kanelakis, Ilias Tevetzidis, Jakub Haluska, Per Bodin, and George Nikolakopoulos. On the design, modeling and experimental verification of a floating satellite platform. *IEEE Robotics and Automation Letters*, 7(2):1364–1371, 2022. doi: 10.1109/LRA.2021.3140134.
- [4] Barış Can Yalçın, Carol Martinez, Sofía Coloma, Ernest Skrzypczyk, and Miguel A Olivares-Mendez. Lightweight floating platform for ground-based emulation of on-orbit scenarios. *IEEE Access*, 11:94575–94588, 2023.
- [5] Pedro Roque, Sujet Phodapol, Elias Krantz, Jaeyoung Lim, Joris Verhagen, Frank J. Jiang, David Dörner, Huina Mao, Gunnar Tibert, Roland Siegwart, Ivan Stenius, Jana Tumova, Christer Fuglesang, and Dimos V. Dimarogonas. Towards open-source and modular space systems with atmos, 2025. URL <https://arxiv.org/abs/2501.16973>.
- [6] Anton Bredenbeck, Shubham Vyas, Willem Suter, Martin Zwick, Dorit Borrmann, Miguel Olivares-Mendez, and Andreas Nüchter. Finding and following optimal trajectories for an overactuated floating robotic platform, 2022. URL <https://arxiv.org/abs/2206.03993>.
- [7] Yashwanth Kumar Nakka, Rebecca C. Foust, Elena Sorina Lupu, David B. Elliott, Irene S. Crowell, Soon-Jo Chung, and Fred Y. Hadaegh. A six degree-of-freedom spacecraft dynamics simulator for formation control research. In *AAS/AIAA Astrodynamics Specialist Conference (AAS 18-476)*, 2018.
- [8] Panagiotis Tsiotras. ASTROS: A 5DoF experimental facility for research in space proximity operations. In *AAS Guidance, Navigation and Control Conference (AAS 14-114)*, 2014.
- [9] Lucas Santaguida and Zheng H. Zhu. Development of air-bearing microgravity testbed for autonomous spacecraft rendezvous and robotic capture control of a free-floating target. *Acta Astronautica*, 203:319–328, 2023.
- [10] Stephen Kwok-Choon, Kathleen Buchala, Briana Blackwell, Steven Lopresti, Markus Wilde, and Thein Htay Go. Design, fabrication, and preliminary testing of air-bearing test vehicles for the study of autonomous satellite maneuvers. In *Proceedings of the 31st Florida Conference on Recent Advances in Robotics (FCRAR)*, page 25, Orlando, FL, USA, May 2018. The paper number in the proceedings is 25.
- [11] Marco Sabatini, Paolo Gasbarri, and Giovanni B. Palmerini. Coordinated control of a space manipulator tested by means of an air bearing free floating platform. *Acta Astronautica*, 139:296–305, 2017.
- [12] Maria G. Bualat, Trey Smith, Ernest E. Smith, Terrence Fong, and D. W. Wheeler. Astrobe: A new tool for ISS operations. In *AIAA SpaceOps Conference*, 2018.
- [13] Trey Smith, Jonathan Barlow, Maria Bualat, Terrence Fong, Christopher Provencher, Hugo Sanchez, and Ernest Smith. Astrobe: A new platform for free-flying robotics on the ISS. In *Int. Symp. Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2016.
- [14] Federico Turchetti, Monica Ekal, Neal Y. Lii, and Maximo A. Roa. Analysis of intra-vehicular robotic free-flyers and their manipulation capabilities. In *75th Int. Astronautical Congress (IAC), Milan*, 2024.
- [15] Matteo El-Hariry, Antoine Richard, Ricard M. Castan, Luis F. W. Batista, Matthieu Geist, Cedric Pradalier, and Miguel Olivares-Mendez. Roboran: A unified robotics framework for reinforcement learning-based autonomous navigation, 2025. URL <https://arxiv.org/abs/2505.14526>.
- [16] Andrej Orsula, Matthieu Geist, Miguel Olivares-Mendez, and Carol Martinez. Space robotics bench: Robot learning beyond earth, 2025. URL <https://arxiv.org/abs/2509.23328>.
- [17] Matteo El-Hariry, Antoine Richard, Vivek Muralidharan, Matthieu Geist, and Miguel Olivares-Mendez. DRIFT: Deep reinforcement learning for intelligent floating platforms trajectories. In *arXiv:2310.04266v2*, 2024.
- [18] Kirk Hovell and Steve Ulrich. On deep reinforcement learning for spacecraft guidance. In *AIAA SciTech Forum*, 2020.
- [19] Dinusha Athauda, Avijit Banerjee, Sumeet Satpute, and George Nikolakopoulos. Intelligent motion planning for collision free autonomous docking of satellite emulation platform using reinforcement learning. *IFAC-*

- PapersOnLine*, 56(2):3354–3359, 2023.
- [20] Yang Cao, Shuang Wang, et al. Reinforcement learning with prior policy guidance for motion planning of dual-arm free-floating space robot. *Aerospace Science and Technology*, 136:108098, 2023.
 - [21] Mayank Mittal et al. Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning. *arXiv preprint arXiv:2511.04831*, 2025. URL <https://arxiv.org/abs/2511.04831>.
 - [22] Jonathan D. Wapman, David C. Sternberg, Kevin Lo, Michael Wang, Laura L. Jones-Wilson, and Swati Mohan. Jet propulsion laboratory small satellite dynamics testbed planar air-bearing propulsion system characterization. *Journal of Spacecraft and Rockets*, 58(4):1126–1140, 2021. doi: 10.2514/1.A34857.
 - [23] Anton Bredenbeck, Shubham Vyas, Willem Suter, Martin Zwick, Dorit Borrmann, Miguel Olivares-Mendez, and Andreas Nüchter. REACSA: Actuated floating platform for orbital robotic concept testing and control software development. In *Proc. ESA Symp. Advanced Space Technologies in Robotics and Automation (ASTRA)*, 2023.
 - [24] Markus Schlotterer and Stephan Theil. Testbed for on-orbit servicing and formation flying dynamics emulation. In *AIAA Guidance, Navigation, and Control Conference*, 2010. doi: 10.2514/6.2010-8108.
 - [25] Youngho Eun, Sang-Young Park, and Geuk-Nam Kim. Development of a hardware-in-the-loop testbed to demonstrate multiple spacecraft operations in proximity. *Acta Astronautica*, 147:48–58, 2018.
 - [26] Zheng Huang, Wei Zhang, Ti Chen, Hao Wen, and Dongping Jin. Characterizing an air-bearing testbed for simulating spacecraft dynamics and control. *Aerospace*, 9(5):246, 2022.
 - [27] Hendrik Kolvenbach and Kjetil Wormnes. Recent developments on ORBIT, a 3-DoF free floating contact dynamics testbed. In *Proc. ESA Symp. Advanced Space Technologies in Robotics and Automation (ASTRA)*, 2015.
 - [28] Evangelos Papadopoulos, Iosif S. Paraskevas, Thaleia Flessa, Kostas Nanos, Georgios Rekleitis, and Ioannis Kontolatis. The NTUA space robot simulator: Design & results. In *Proc. ESA Symp. Advanced Space Technologies in Robotics and Automation (ASTRA)*, 2015.
 - [29] James Ragan, Benjamin Riviere, Fred Y. Hadaegh, and Soon-Jo Chung. Online tree-based planning for active spacecraft fault estimation and collision avoidance. *Science Robotics*, 2024.
 - [30] Franek Stark, Shubham Vyas, Georg Schildbach, and Frank Kirchner. Linear model predictive control for a planar free-floating platform: A comparison of binary input constraint formulations, 2023. URL <https://arxiv.org/abs/2312.10788>.
 - [31] Angel Flores-Abad, Ou Ma, Khanh Pham, and Steve Ulrich. A review of space robotics technologies for on-orbit servicing. *Progress in Aerospace Sciences*, 68:1–26, 2014. doi: 10.1016/j.paerosci.2014.03.002.
 - [32] S. Ali A. Moosavian and Evangelos Papadopoulos. Free-flying robots in space: An overview of dynamics modeling, planning and control. *Robotica*, 25(5):537–547, 2007. doi: 10.1017/S0263574707003438.
 - [33] Yoji Umetani and Kazuya Yoshida. Resolved motion rate control of space manipulators with generalized Jacobian matrix. *IEEE Transactions on Robotics and Automation*, 5(3):303–314, 1989.
 - [34] Evangelos Papadopoulos and Steven Dubowsky. On the nature of control algorithms for free-floating space manipulators. *IEEE Transactions on Robotics and Automation*, 7(6):750–758, 1991.
 - [35] Kazuya Yoshida. Engineering test satellite VII flight experiments for space robot dynamics and control: Theories on laboratory test beds ten years ago, now in orbit. *The International Journal of Robotics Research*, 22(5):321–335, 2003. doi: 10.1177/0278364903022005003.
 - [36] Andrew Ogilvie, Justin Allport, Michael Hannah, and John Lymer. Autonomous satellite servicing using the orbital express demonstration manipulator system. In *Proc. 9th Int. Symp. Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2008.
 - [37] Jordi Artigas, Marco De Stefano, Wolfgang Rackl, Roberto Lampariello, Bernhard Brunner, Wieland Bertl-eff, Robert Burger, Oliver Porges, Alessandro Massimo Giordano, Christoph Borst, and Alin Albu-Schäffer. The OOS-SIM: An on-ground simulation facility for on-orbit servicing robotic operations. In *IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 2854–2860, 2015. doi: 10.1109/ICRA.2015.7139588.
 - [38] Tomasz Rybus, Karol Seweryn, Jakub Oles, Fatina L. Basmadji, Kamil Tarenko, Rafal Moczydlowski, Tomasz Barcinski, Jan Kindracki, Lukasz Mezyk, Pawel Paszkiewicz, and Piotr Wolanski. Application of a planar air-bearing microgravity simulator for demonstration of operations required for an orbital capture with a manipulator. *Acta Astronautica*, 155:211–229, 2019. doi: 10.1016/j.actaastro.2018.12.004.
 - [39] Kenji Nagaoka, Ryota Kameoka, and Kazuya Yoshida. Repeated impact-based capture of a spinning object by a dual-arm space robot. *Frontiers in Robotics and AI*, 5: 115, 2018. doi: 10.3389/frobt.2018.00115.
 - [40] Brian Gaudet, Richard Linares, and Roberto Furfaro. Deep reinforcement learning for six degree-of-freedom planetary landing. *Advances in Space Research*, 65(7): 1723–1741, 2020.
 - [41] Brian Gaudet, Richard Linares, and Roberto Furfaro. Terminal adaptive guidance via reinforcement meta-learning: Applications to autonomous asteroid close-proximity operations. *Acta Astronautica*, 171:1–13, 2020. doi: 10.1016/j.actaastro.2020.02.036.
 - [42] Wiebke Retagne, Jonas Dauer, and Günther Waxenegger-Wilfing. Adaptive satellite attitude control for varying masses using deep reinforcement learning. *Frontiers in Robotics and AI*, 11:1402846, 2024. doi: 10.3389/frobt.2024.1402846.
 - [43] Kirk Hovell and Steve Ulrich. Deep reinforcement learning for spacecraft proximity operations guidance. *Journal of Spacecraft and Rockets*, 58(2):254–264, 2021. doi: 10.2514/1.A34838.
 - [44] Kirk Hovell and Steve Ulrich. Laboratory experi-

- mentation of spacecraft robotic capture using deep-reinforcement-learning-based guidance. *Journal of Guidance, Control, and Dynamics*, 45(11):2138–2146, 2022. doi: 10.2514/1.G006656.
- [45] Matteo El Hariry, Andrea Cini, Giacomo Mellone, and Alessandro Balossino. Deep reinforcement learning policies for underactuated satellite attitude control, 2025. URL <https://arxiv.org/abs/2505.00165>.
- [46] Samantha Chapin, Kenneth Stewart, Roxana Leontie, and Carl Glen Henshaw. Autonomous planning in-space assembly reinforcement-learning free-flyer (APIARY): International space station astrobbee testing, 2025. URL <https://arxiv.org/abs/2512.03729>.
- [47] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance gpu-based physics simulation for robot learning. In *Proc. NeurIPS Track on Datasets and Benchmarks*, 2021. URL <https://arxiv.org/abs/2108.10470>.
- [48] Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Proc. 5th Conf. on Robot Learning (CoRL)*, volume 164 of *Proceedings of Machine Learning Research*, pages 91–100, 2022. URL <https://proceedings.mlr.press/v164/rudin22a.html>.
- [49] Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872, 2019. doi: 10.1126/scirobotics.aau5872.
- [50] Qiayuan Liao, Bike Zhang, Xuanyu Huang, Xiaoyu Huang, Zhongyu Li, and Koushil Sreenath. Berkeley Humanoid: A research platform for learning-based control. In *arXiv preprint*, 2024.
- [51] Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Champion-level drone racing using deep reinforcement learning. *Nature*, 620(7976):982–987, 2023. doi: 10.1038/s41586-023-06419-4.
- [52] David B. D’Ambrosio et al. Robotic table tennis: A case study into a high speed learning system. In *Robotics: Science and Systems (RSS)*, 2023.
- [53] Ruben Grandia, Espen Knoop, Michael A. Hopkins, Georg Wiedebach, Jared Bishop, Steven Pickles, David Müller, and Moritz Bächer. Design and control of a bipedal robotic character. In *Robotics: Science and Systems (RSS)*, 2024.
- [54] New Way Air Bearings. 150mm flat round air bearing (sku: S1015001). <https://www.newwayairbearings.com/catalog/product/150mm-flat-round-air-bearings/>, 2026.
- [55] ODrive Robotics. Odrive: High-performance motor control. <https://odriverobotics.com/>, 2026. Accessed: July 20, 2026.
- [56] Clemens Schwarke, Mayank Mittal, Nikita Rudin, David Hoeller, and Marco Hutter. Rsl-rl: A learning library for robotics research. *arXiv preprint arXiv:2509.10771*, 2025.
- [57] Mayank Mittal, Calvin Yu, Qinxi Yu, Jingzhou Liu, Nikita Rudin, David Hoeller, Jia Lin Yuan, Ritvik Singh, Yunrong Guo, Hammad Mazhar, Ajay Mandlekar, Buck Babich, Gavriel State, Marco Hutter, and Animesh Garg. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023. doi: 10.1109/LRA.2023.3270034.
- [58] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL <https://arxiv.org/abs/1707.06347>.
- [59] Rudolf E. Kálmán. Contributions to the theory of optimal control. *Boletín de la Sociedad Matemática Mexicana*, 5(1):102–119, 1960.

APPENDIX

A. Reproducibility and Open Access

To support reproducibility and lower the barrier to entry for learning-based spacecraft GNC research, we openly release all assets associated with this work. The complete hardware design (CAD and BOM), firmware, ROS 2 packages, simulation environments, training scripts, experimental logs, and supplementary videos are available at <https://snt-spacer.github.io/UniLuFP/>.

B. More Implementation Details

All policies are trained with Proximal Policy Optimization (PPO), using the RSL-RL [56] implementation within IsaacLab [21], across 5 random seeds, with 4096 parallel environments during training and 640 environments at evaluation. We consider two policy architectures: a feed-forward multi-layer perceptron (MLP) and a recurrent variant that inserts a single-layer gated recurrent unit (GRU) before the actor and critic heads to provide memory for rejecting the unobserved, time-varying disturbances induced by arm motion. The two architectures share all optimization hyperparameters; they differ only in the recurrent core and in the training budget (the recurrent policy is trained for twice as many iterations to accommodate the additional parameters). Table V summarizes both configurations.

a) *Domain randomization*: To bridge the sim-to-real gap, a set of per-environment physical parameters of the platform base is re-sampled at every episode reset. The *base mass* is perturbed by a uniform offset in $[-5, +5]$ kg. The *center of mass* is displaced by up to 0.1 m in a uniformly random horizontal direction. A *constant external wrench* is drawn once per episode and held for its entire duration: a planar bias force of up to 1 N along a random direction, together with a bias yaw torque of up to 0.05 N m, emulating unmodeled effects such as air-bearing tilt, residual table slope, and thruster imbalance. These three are the randomizations enabled during training; action-level noise and thruster-gain rescaling are implemented in the framework but disabled for the reported experiments. The platform’s moment of inertia is not randomized explicitly; inertial variation is instead induced physically through the three static LEVION arm configurations (Fig. 12) and implicitly through the mass and CoM randomization.

TABLE V: PPO hyperparameters for the Dynamic Disturbance Rejection policies, for the feed-forward (MLP) and recurrent (GRU) variants. Shared values span both columns.

Hyperparameter	PPO (MLP)	PPO (GRU)
Number of envs		4096
Number of steps per iteration		16
Number of learning epochs		5
Number of mini-batches		4
Clip range		0.2
Entropy coefficient		0.005
Value loss coefficient		1.0
GAE balancing factor λ		0.95
Discount factor γ		0.99
Desired KL-divergence		0.01
Initial learning rate		1×10^{-3}
Learning-rate schedule		Adaptive
Max gradient norm		1.0
Initial action noise std		1.0
Actor and critic MLP		[64, 64]
Activation function		Tanh
Recurrent core	—	GRU
GRU hidden dimension	—	64
GRU layers	—	1
Training iterations	1000	2000

Table VII details the reward terms used to train each policy. Every reward is the weighted sum of a task objective and a set of regularization terms that shape *how* the objective is achieved. For **Point-to-Pose Navigation** (a), the dominant term is an exponential pose reward that rewards simultaneous convergence in position (d_p) and heading (d_h); this is complemented by soft linear- and angular-velocity penalties that discourage overshoot, a boundary term that keeps the platform inside the workspace, and the platform regularization of block (c). The latter is a set of small penalties on the actuation itself—joint acceleration, thruster action rate and effort, reaction-wheel speed and torque, and arm motion and collisions—that together suppress chattering, discourage wasteful actuation, and yield smooth commands. As noted in the table, the reaction-wheel and arm penalties are active only for the actuator subsets that include those effectors.

The **Reaction-Wheel Stabilization** task (b) is the momentum-damping experiment: the platform spawns with a residual body spin of 0.1–0.5 rad/s and must null its yaw rate using only the reaction wheel. The single task term rewards driving the yaw rate ω_z toward zero ($\omega_z^* = 0$) through a sharp exponential kernel ($\lambda_\omega = 0.5$), so the reward grows steeply as the platform comes to rest. Because this policy actuates only the reaction wheel, its regularization is wheel-specific rather than the terms in (c): a quadratic penalty on internal wheel speed ω_{rw} discourages saturation (a saturated wheel produces no usable torque), and a penalty on commanded torque τ_{rw} discourages unnecessary actuation, encouraging the policy to brake decisively and then desaturate. These two penalties are deliberately small relative to the unit-weight stabilization term, so they regularize the behavior without competing with the primary objective.

C. More Experiments

Tables VI and VIII report the complete ablation study underlying the curated hardware numbers of Section IV. Across

Method	e_p [m] ↓	e_o [rad] ↓
(a) Point-to-Pose Navigation — all actuators		
PPO (Sim)	0.0022 (± 0.002)	0.0044 (± 0.003)
PPO + DR (Sim)	0.0274 (± 0.076)	0.0192 (± 0.019)
PPO + DR (Real)	0.0584 (± 0.041)	0.0560 (± 0.015)
PPO-GRU + DR, 1000 ep (Sim)	0.0209 (± 0.034)	0.0277 (± 0.026)
PPO-GRU + DR, 1000 ep (Real, chkpt. 1)	0.0862 (± 0.086)	0.0443 (± 0.031)
PPO-GRU + DR, 1000 ep (Real, chkpt. 2)	0.0855 (± 0.016)	0.0211 (± 0.017)
PPO-GRU + DR, 2000 ep (Sim)	0.0124 (± 0.019)	0.0129 (± 0.012)
PPO-GRU + DR, 2000 ep (Real, chkpt. 1)	0.0214 (± 0.048)	0.0484 (± 0.024)
PPO-GRU + DR, 2000 ep (Real, chkpt. 2)	0.0081 (± 0.001)	0.0257 (± 0.015)
(b) Dynamic Disturbance Rejection		
PPO (Sim)	0.012 (± 0.002)	0.016 (± 0.003)
PPO + DR (Sim)	0.029 (± 0.003)	0.033 (± 0.004)
PPO + DR (Real, seed 1)	0.1346 (± 0.041)	0.0438 (± 0.036)
PPO + DR (Real, seed 3)	0.1502 (± 0.029)	0.0708 (± 0.018)
PPO-GRU (Sim)	0.0070 (± 0.005)	0.0074 (± 0.011)
PPO-GRU + DR (Sim)	0.0156 (± 0.016)	0.0340 (± 0.057)
PPO-GRU + DR (Real, seed 4)	0.1521 (± 0.044)	0.0887 (± 0.074)
PPO-GRU + DR (Real, seed 1)	0.0889 (± 0.041)	0.0346 (± 0.018)

TABLE VI: **All-actuator navigation and dynamic disturbance rejection — full ablation.** Final position error e_p and heading error e_o for the two studies in which the arms are actuated (all-actuators navigation) or continuously animated as a disturbance (dynamic rejection), so no static arm pose applies. For the all-actuators policy we also sweep the recurrent training budget (1000 vs. 2000 epochs); for the hardware entries we report multiple training seeds / deployment runs. Lower is better.

tasks we sweep up to four axes. **Arm configuration** varies the platform’s effective moment of inertia by fixing the LEVION arms at one of the three static poses of Fig. 12: *Side*, *Rest*, and *Closed*. This axis applies only to the reduced-actuator subsets (thrusters only, thrusters + reaction wheel); the all-actuators navigation policy and the disturbance-rejection policy either control or continuously animate the arms and therefore have no single static pose. **Policy architecture** contrasts the feed-forward PPO (MLP) policy with the recurrent PPO-GRU policy, whose hidden state provides the memory needed to infer unobserved inertial shifts. **Domain randomization** (DR) toggles randomization of the base mass, center of mass, and a constant external wrench (a per-episode bias force and torque) during training, and is a prerequisite for hardware transfer; the exact ranges are given in Appendix B. **Deployment** distinguishes simulation (Sim) from hardware (Real); for the hardware entries we additionally report several independent training seeds / deployment runs to expose the run-to-run spread of sim-to-real transfer.

Two trends emerge from Table VIII. First, in simulation both architectures reach sub-centimeter position error in every arm configuration; adding DR raises the simulation error only marginally while enabling hardware transfer, the expected robustness–accuracy trade-off. Second, the arm configuration has a comparatively small effect on the learned controllers relative to the sim-to-real gap, confirming that the policies absorb the inertia change induced by the static arm pose rather than overfitting to a particular one. On hardware, the feed-forward PPO+DR policy transfers on par with or slightly better than PPO-GRU+DR for this *static* navigation task,

TABLE VII: Reward terms used to train the learned control policies. Each block lists a task-specific reward, given as the weighted sum of its rows. Block (c) collects the platform-regularization penalties added to the *Point-to-Pose Navigation* policy (a); as noted per row, individual terms apply only to configurations that include the corresponding actuators (reaction wheel, arms). The reaction-wheel *Stabilization* policy (b) actuates only the wheel and is therefore regularized by just the reaction-wheel saturation and usage terms, which share the definitions and weights of the matching rows in (c). Distances $d_\bullet$ denote the absolute error between the current and target quantity; negative weights indicate penalties.

Term	Expression	Weight	Scalar	Description
(a) Point-to-Pose Navigation				
Pose error	$r_t^{\text{point-pose-max}} = r_t^{\text{pose}} + r_t^{\text{vel}} + r_t^{\text{head}}$			
Linear Velocity	$\alpha_{\text{pose}}(\exp(-d_p/\lambda_p) \times \exp(-d_h/\lambda_h))$	$\alpha_{\text{pose}} = 1.0$	$\lambda_p, \lambda_h = 1.0$	Precise spatial convergence and alignment to target.
Angular Velocity	$\alpha_v \text{clip}(v - v_{\min}, 0, v_{\max} - v_{\min})$	$\alpha_v = -0.05$	$v_{\min} = 0.5, v_{\max} = 2.0$	Velocity regulation and movement encouragement.
Boundary	$\alpha_{\omega} \text{clip}(\omega - \omega_{\min}, 0, \omega_{\max} - \omega_{\min})$	$\alpha_{\omega} = -0.05$	$\omega_{\min} = 0.5, \omega_{\max} = 20.0$	Control of rotational stability.
	$\alpha_{\text{bad}} \exp(-d_b/\lambda_b)$	$\alpha_{\text{bad}} = -10.0$	$\lambda_b = 1.0$	Safety constraint to stay within workspace.
(b) Reaction-Wheel Stabilization (Momentum Damping)				
Angular-rate nulling	$\alpha_{\omega} \exp(- \omega_z - \omega_z^* /\lambda_{\omega})$	$\alpha_{\omega} = 1.0$	$\lambda_{\omega} = 0.5, \omega_z^* = 0$	Rewards driving the platform yaw rate ω_z to zero, i.e. dissipating the initial angular momentum.
Reaction-wheel saturation	$\gamma_{\text{sat}} \frac{\omega_{\text{rw}}}{\omega_{\text{thr}}}$	$\gamma_{\text{sat}} = -2.5 \times 10^{-6}$		Penalizes high internal wheel speed ω_{rw} to keep the wheel from saturating and preserve control authority.
Reaction-wheel usage	$\gamma_{\text{use}} \tau_{\text{rw}} $	$\gamma_{\text{use}} = -0.05$		Penalizes the commanded wheel torque τ_{rw} to discourage unnecessary actuation.
(c) Platform Regularization (Navigation)				
Joint acceleration	$r_t^{\text{joint-acc}} = r_t^{\text{acc}} + r_t^{\text{rate}} + r_t^{\text{eff}} + r_t^{\text{sat}} + r_t^{\text{use}} + r_t^{\text{arm}} + r_t^{\text{coll}}$			
Thruster action rate	$\gamma_{\text{acc}} \sum_i \ddot{q}_i^2$	$\gamma_{\text{acc}} = -2.5 \times 10^{-6}$		Penalizes summed squared joint accelerations for smoother motion. (<i>all configs</i>)
Thruster effort	$\gamma_{\text{rate}} \sum_i \dot{a}_i - \dot{a}_{i-1} $	$\gamma_{\text{rate}} = -0.015$		Penalizes step-to-step change in thruster commands to reduce chattering. (<i>all configs</i>)
Reaction-wheel saturation	$\gamma_{\text{eff}} \sum_i f_i /f_{\text{max}}$	$\gamma_{\text{eff}} = -0.01$	$f_{\text{max}} = 1.0$	Penalizes total normalized thrust to reduce sustained actuation. (<i>all configs</i>)
Reaction-wheel usage	$\gamma_{\text{sat}} \frac{\omega_{\text{rw}}}{\omega_{\text{thr}}}$	$\gamma_{\text{sat}} = -2.5 \times 10^{-6}$		Penalizes high wheel speed to avoid saturation. (<i>configs with reaction wheel</i>)
Arm action rate	$\gamma_{\text{use}} \tau_{\text{rw}} $	$\gamma_{\text{use}} = -0.05$		Penalizes commanded wheel torque. (<i>configs with reaction wheel</i>)
Arm collision	$\gamma_{\text{arm}} \sum_i \theta_i - \theta_{i-1} $	$\gamma_{\text{arm}} = -0.1$		Penalizes step-to-step change in arm position targets to limit induced vibration. (<i>all-actuators config</i>)
	$\gamma_{\text{coll}} \mathbb{1}[F_{\text{arm}} > F_{\text{thr}}]$	$\gamma_{\text{coll}} = -1.0$	$F_{\text{thr}} = 0.1 \text{ N}$	Penalizes contact detected on the arm links. (<i>all-actuators config</i>)

TABLE VIII: **Point-to-Pose Navigation - static-arm ablation.** Final position error e_p and heading error e_o for the reduced-actuator subsets, evaluated across the three static LEVION arm poses (*Side, Rest, Closed*) of Fig. 12. Each subset is ablated over policy architecture (feed-forward PPO vs. recurrent PPO-GRU), domain randomization (DR), and deployment (Sim vs. Real hardware). Entries are mean $\pm$ std over evaluation episodes; lower is better.

Method	Side		Rest		Closed	
	e_p [m] $\downarrow$	e_o [rad] $\downarrow$	e_p [m] $\downarrow$	e_o [rad] $\downarrow$	e_p [m] $\downarrow$	e_o [rad] $\downarrow$
(a) Thrusters and reaction wheel						
PPO (Sim)	0.0021 (± 0.001)	0.0038 (± 0.004)	0.0031 (± 0.003)	0.0020 (± 0.002)	0.0034 (± 0.002)	0.0018 (± 0.001)
PPO + DR (Sim)	0.0155 (± 0.042)	0.0193 (± 0.018)	0.0204 (± 0.062)	0.0406 (± 0.051)	0.0154 (± 0.036)	0.0122 (± 0.012)
PPO + DR (Real)	0.0115 (± 0.008)	0.0302 (± 0.009)	0.0246 (± 0.016)	0.0313 (± 0.014)	0.0406 (± 0.023)	0.0121 (± 0.007)
PPO-GRU (Sim)	0.0137 (± 0.011)	0.0047 (± 0.004)	0.0220 (± 0.022)	0.0047 (± 0.004)	0.0045 (± 0.006)	0.0040 (± 0.003)
PPO-GRU + DR (Sim)	0.0083 (± 0.009)	0.0081 (± 0.009)	0.0097 (± 0.010)	0.0069 (± 0.006)	0.0145 (± 0.011)	0.0067 (± 0.007)
PPO-GRU + DR (Real)	0.0547 (± 0.042)	0.0275 (± 0.015)	0.0482 (± 0.027)	0.0279 (± 0.020)	0.0679 (± 0.050)	0.0598 (± 0.035)
(b) Thrusters only						
PPO (Sim)	0.0025 (± 0.002)	0.0028 (± 0.002)	0.0047 (± 0.005)	0.0028 (± 0.004)	0.0047 (± 0.005)	0.0027 (± 0.002)
PPO + DR (Sim)	0.0080 (± 0.014)	0.0086 (± 0.007)	0.0189 (± 0.073)	0.0104 (± 0.017)	0.0124 (± 0.045)	0.0087 (± 0.009)
PPO + DR (Real)	0.0087 (± 0.007)	0.0202 (± 0.013)	0.0176 (± 0.020)	0.0160 (± 0.006)	0.0348 (± 0.013)	0.0310 (± 0.041)
PPO-GRU (Sim)	0.0064 (± 0.011)	0.0024 (± 0.003)	0.0057 (± 0.009)	0.0046 (± 0.003)	0.0099 (± 0.011)	0.0055 (± 0.004)
PPO-GRU + DR (Sim)	0.0069 (± 0.009)	0.0087 (± 0.008)	0.0056 (± 0.007)	0.0105 (± 0.008)	0.0075 (± 0.009)	0.0062 (± 0.005)
PPO-GRU + DR (Real)	0.0241 (± 0.016)	0.0335 (± 0.016)	0.0745 (± 0.040)	0.0219 (± 0.003)	0.0527 (± 0.019)	0.0414 (± 0.036)

consistent with the recurrent policy being most valuable when the disturbance is time-varying (Table VI) rather than fixed.

For the all-actuators and disturbance-rejection studies (Table VI), the recurrent training-length sweep shows that the longer 2000-epoch schedule yields the best-transferring all-actuators policy ($e_p = 0.0081 \text{ m}$ on hardware). Under the continuous CoM disturbance the recurrent policy clearly outperforms the feed-forward one in simulation ($e_p = 0.0070$ vs. 0.012 m), and its best hardware seed attains the lowest real-world position error ($e_p = 0.0889 \text{ m}$), motivating the recurrent policy for the dynamic task. The larger hardware position errors in this task, relative to static navigation, reflect the never-settling nature of the injected disturbance discussed in Section IV.