POLITECNICO DI TORINO
Repository ISTITUZIONALE

MPC for underactuated spacecraft control with a Lyapunov supervised physics-informed neural network correction layer

Original
MPC for underactuated spacecraft control with a Lyapunov supervised physics-informed neural network correction layer / Ayanmanesh Motlaghmofrad, A., Cena, C., Martini, M., Chiaberge, M.. - ELETTRONICO. - (2026), pp. 81-85. (The 3rd Conference on AI in and for Space (SPAICE 2026) Noordwijk (Ned) 21-23 October 2026) [10.5281/zenodo.23495602].

Availability:
This version is available at: 11583/3011893 since: 2026-10-27T16:16:17Z

Publisher:
Zenodo

Published
DOI:10.5281/zenodo.23495602

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

12 November 2026

MPC for underactuated spacecraft control with a Lyapunov supervised physics-informed neural network correction layer

Amirhossein Ayanmanesh Motlaghmofrad[1], Carlo Cena[*†1,2], Mauro Martini[1], Marcello Chiaberge[1]

[1] *Politecnico di Torino, Torino (TO), ITA*
[2] *Argotec S.R.L., San Mauro (TO), ITA*

Underactuated spacecraft faces controllability limitations and heightened sensitivity to environmental disturbances, complicating attitude maneuvering and stabilization. Due to the lack of control authority along the underactuated axis, conventional controllers cannot directly stabilize all attitude components and therefore require reference planning strategies. Furthermore, MPC approaches remain sensitive to inertia uncertainty and unmodeled dynamic couplings, resulting in degraded tracking performance under mismatch. To address these issues, we consider a hierarchical architecture integrating three layers: (i) a nonlinear model predictive controller (NMPC) for constraint and underactuation-aware maneuver planning and nominal closed-loop stability under actuator limits; (ii) a physics-informed neural network (PINN) trained offline on simulation data to estimate residual disturbance torques, with loss terms that enforce consistency with rigid-body rotational dynamics; (iii) a Lyapunov-based supervisory safety mechanism that evaluates the learned correction online and bounds or suppresses its influence to preserve the stability properties of the baseline controller. The architecture is evaluated in a high-fidelity simulation environment modelling reaction wheel dynamics, actuator saturation, and environmental disturbances. Experimental studies show statistically significant reductions in steady-state attitude error relative to standalone NMPC while maintaining robust behavior under uncertainty. The supervisory layer ensures graceful degradation to purely model-based control when the learning-based augmentation is unreliable.

1 Introduction

Attitude control is fundamental to spacecraft operations, enabling precise maneuvers, Earth and deep-space observation, and reliable communication between space and ground. However, design decisions or actuator failures can result in underactuated spacecraft, reducing control over rotational dynamics and necessitating careful characterization of the actuation system. Common actuators include control moment gyros (CMGs), reaction wheels (RWs), and thrusters. When only internal torques are available, control is constrained by the conservation of angular momentum, leading to fundamental accessibility limitations [1]. For the study of underactuated spacecraft, two principal approaches have been proposed. The first assumes zero total angular momentum, which simplifies the dynamics and enables short-time local controllability of the reduced attitude system, allowing rest-to-rest maneuvers between admissible configurations [2]. The second exploits predictable external torques, such as solar radiation pressure, to recover full three-axis attitude control without relying on conservation assumptions [3, 4]. With a practical approach in mind, the present study adopts the zero-total-angular-momentum assumption. Although environmental disturbances may induce drift in the otherwise uncontrollable component of angular momentum over time, it is assumed that occasional external momentum management counteracts this drift. Recent advancements in space technology have explored neural networks for trajectory optimization and adaptive feedback control, suitable for onboard implementation [5, 6]. Applications include adaptive neural-network controllers [7], neural networks assisting primary controllers [8], and neural-network-based optimal attitude control with impulsive thrusters [9]. Physics-informed approaches, such as physics-informed normalizing flows, have also been used in several applications [10–13]. Additionally, in the broader control field data-driven MPC approaches are increasingly explored [6, 14]. This work proposes a hybrid model-based and learning-based control architecture that systematically combines the strengths of both paradigms while mitigating their standalone limitations. A model-based baseline controller ensures nominal stability, for the controllable axes, and safety properties, forming a certifiable backbone suitable for safety-critical operations. This baseline is augmented with a physics-informed neural network (PINN) operating as a residual torque compensator, enabling improved tracking performance in the presence of underactuation and structured modeling uncertainties. To preserve closed-loop guarantees, a deterministic Lyapunov-based supervisory mechanism constrains the neural network contribution in real time, ensuring that stability and safety properties remain intact in accordance with formal guidance [15].

The main contributions of this paper can be summarized as follows:

[*]Corresponding author. E-Mail: carlo.cena@polito.it
[†]This work was supported by the the project PNRR-NGEU which has received funding from the MUR – DM 117/2023.

3rd Conference on AI in and for Space (SPAICE 2026)
© 2026 A. Ayanmanesh Motlaghmofrad et al., licenced via CC BY 4.0.

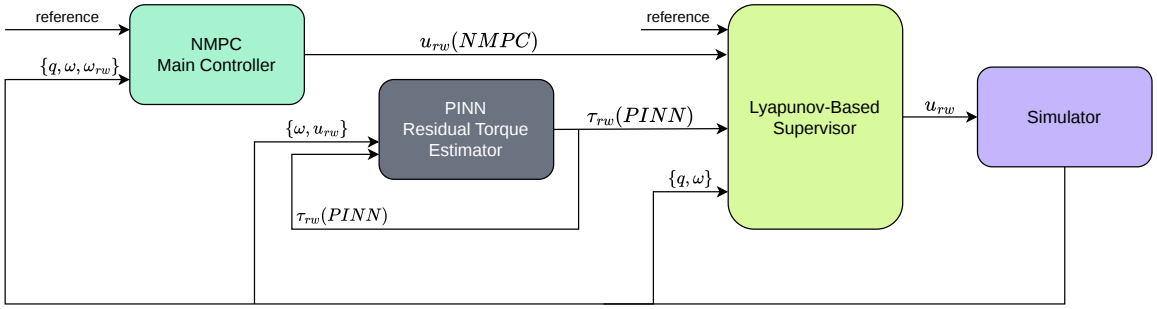


Figure 1: High-level block diagram representation of the proposed control system architecture for the underactuated spacecraft under study.

- We propose a novel hybrid control framework that combines a baseline model-based controller with a physics-informed neural network (PINN) to address underactuation and modeling uncertainties, while preserving theoretical guarantees for safety-critical operations.
- We develop a deterministic Lyapunov-based supervisory mechanism that constrains the neural network contribution in real time, in accordance with formal verification principles [15].
- We show improved steady-state tracking performance, achieving a 3.8% reduction in average steady-state error and a 24.16% reduction for the worst case compared to the baseline controller, and demonstrate significant improvements in the final error, yielding a 12.7% reduction in average final error, 52.9% reduction in the maximum final error, and 29.42% reduction in final error variance.

2 Methodology

The control architecture is hierarchical and comprises three main components: a nonlinear model predictive controller (NMPC) as the baseline policy, a physics-informed neural network (PINN) that estimates residual disturbance torques to reduce model-plant mismatch, and a Lyapunov-based supervisory safety layer that ensures stability and constraint satisfaction. The attitude is represented by a unit quaternion q and body rates ω. The reaction-wheel speeds are denoted by ω_{RW}, and the input is the wheel motor torque u_{RW}. Practical limits are treated as hard constraints, including wheel speed $|\omega_{\mathrm{RW}}| \leq 6000\,\mathrm{rpm}$, motor torque $\|u_{\mathrm{RW}}\|_\infty \leq 0.05\,\mathrm{Nm}$, stored momentum $\approx 0.9\,\mathrm{Nms}$, and torque slew-rate $\|\Delta u_{\mathrm{RW}}\|_\infty \leq \Delta u_{\max}$ with $\dot{u}_{\max} = 0.4\,\mathrm{Nm/s}$ and $T_s = 0.5\,\mathrm{s}$. The nominal predictive model assumes a diagonal inertia J with underactuated actuator distribution and discretization by fixed-step RK4. Gravity-gradient and aerodynamic drag are included, while solar radiation pressure and magnetic torques are neglected as secondary in the considered regime. The control objective is to achieve a steady-state pointing error bounded by a 0.5° settling-angle requirement, satisfy all actuator and operational constraints at all times, and preserve robust closed-loop behavior under bounded model mismatch. The NMPC formulation employs a discrete prediction model from RK4 with an update period $T_{\mathrm{NMPC}} = 0.5\,\mathrm{s}$ and horizon length $N_h = 24$ (12s window). The optimal control problem minimizes a finite-horizon cost with stage terms penalizing attitude error, angular-rate error, control effort, and input increments, plus a quadratic terminal cost. Whereas, the PINN learns a bounded residual torque τ_d as a corrective term in the augmented discrete-time dynamics of Equation (1).

$$x_{k+1} = \Phi_{\mathrm{RK4}}(x_k, u_{\mathrm{RW},k}, \tau_d(k)), \tag{1}$$

where $\Phi_{\mathrm{RK4}}(\cdot)$ denotes the fourth-order Runge–Kutta discretization of the continuous dynamics. The residual torque is assumed constant over each NMPC sampling interval, $\tau_d(t) \approx \tau_d(k)$ for $t \in [kT_s, (k+1)T_s)$, reflecting the fact that dominant model mismatch sources, such as inertial parameter errors and slowly varying disturbances, evolve on time scales significantly slower than the NMPC update rate. The network follows a NARX structure with three hidden layers of 128 neurons each with tanh activations. It receives as input 16 features: the previous 2 estimated disturbance torque, and a set of 14 engineered features created starting from the current and previous spacecraft angular velocity, and the current RWs torques. Inputs are normalized using training set statistics, and outputs are scaled by conservative bounds to enforce magnitude limits. The training objective is formulated as a multi-step rollout loss through the dynamics. Starting from state x_k, the augmented dynamics are propagated over $H = 10$ integration steps (sampled ten times faster than the NMPC update frequency), yielding a predicted angular velocity $\hat{\omega}_{k+H}(\theta)$. The physics-informed loss is defined in Equation (2).

$$\mathcal{L}(\theta) = \sum_{k \in \mathcal{K}} \|\hat{\omega}_{k+H}(\theta) - \omega_{k+H}\|_2^2, \tag{2}$$

where ω_{k+H} denotes the measured angular velocity from simulation data, $\mathcal{K}$ is the set of training time indices, and θ

© 2026 A. Ayanmanesh Motlaghmofrad et al., licenced via CC BY 4.0.

represents the network parameters. By construction, the loss depends on θ only through the effect of $\tau_d(k)$ on state evolution over the prediction horizon. This formulation aligns the learning objective with the intended role of the PINN: improving short-horizon prediction accuracy in a manner consistent with its deployment in the closed-loop system. The Lyapunov-based supervisory layer does not stabilize the spacecraft independently; its role is to avoid that the learned correction degrades the closed-loop stability already achieved by the nominal NMPC. At each control update, the supervisor evaluates whether injecting the PINN correction is compatible with a non-increasing Lyapunov-like energy associated with the attitude error dynamics; if this condition is violated, the correction is attenuated or rejected while the nominal NMPC command is left unchanged.

Let $q_e = q^{-1} \otimes q_{\text{ref}}$ and $\omega_e = \omega - \omega_{\text{ref}}$ denote the attitude and rate tracking errors. The candidate Lyapunov function is

$$V(q_e, \omega_e) = \tfrac{1}{2}\omega_e^\top J \omega_e + k_q\,(1 - q_{e,0}), \qquad k_q > 0, \quad (3)$$

positive definite about $q_e = [1, 0, 0, 0]^\top$, $\omega_e = 0$. For the rest-to-rest objective considered here ($\dot{\omega}_{\text{ref}} \approx 0$), its derivative along the rigid-body dynamics $J\dot{\omega} = \tau - \omega \times (J\omega)$ is

$$\dot{V} = \omega_e^\top \big(\tau - \omega \times (J\omega)\big) - \frac{k_q}{2} q_{e,0}\, q_{e,v}^\top \omega_e. \quad (4)$$

The applied torque is decomposed as $\tau = \tau_{\text{nom}} + \tau_{\text{corr}}$, where $\tau_{\text{nom}} = -B_{\text{act}} u_{\text{NMPC}}$ is generated by the NMPC and $\tau_{\text{corr}} = -B_{\text{act}} u_{\text{PINN}}$ is the realization of the PINN's disturbance estimate τ_d. Since an ideal correction satisfies $\tau_{\text{corr}} = -\tau_d$, a learning failure with a sign-flipped estimate would double rather than cancel the disturbance term in Equation (1). To guard against this failure mode, the supervisor tests the Lyapunov condition under a conservative assumption: it introduces a scalar attenuation factor $\alpha \in [0, 1]$ and evaluates $\dot{V}$ under the test torque

$$\tau_{\text{test}} = \tau_{\text{nom}} - 2\alpha\, \tau_{\text{corr}}, \quad (5)$$

selecting the largest α for which $\dot{V}(\tau_{\text{test}}) \le 0$; if no $\alpha > 0$ satisfies this condition, the correction is rejected ($\alpha = 0$) and the closed loop reverts to the pure NMPC command. This condition is enforced as a local safeguard. Additional conservative bounds are applied at implementation level: the PINN correction torque is saturated to $0.015\,\text{Nm}$, the total commanded torque remains within the $0.05\,\text{Nm}$ actuator limit, and the correction is disabled whenever the attitude error falls below 1°. Overall, the closed loop consists of the NMPC baseline performing both implicit trajectory planning and feedback control, to generate constraint-aware command references that actively reduce the attitude error along the underactuated axis, while ensuring nominal stability and recursive feasibility; the PINN providing a bounded, physics-regularized residual torque estimate that drives the spacecraft's behavior closer to the NMPC's predictive model; and the Lyapunov supervisor enforcing a safety cage that guarantees graceful degradation to pure NMPC when learning is unreliable.

3 Experimental Setup

The validation campaign quantitatively evaluates whether PINN-based residual torque compensation improves attitude regulation performance relative to a baseline NMPC controller. All experiments are conducted under controlled and repeatable simulation conditions with fixed controller parameters established prior to testing. We performed a testing campaign spanning diverse initial attitudes, by generating 50 unit direction vectors uniformly distributed on the unit sphere $\mathbb{S}^2$ using a Fibonacci lattice sampling scheme. For each direction, two initial attitudes are defined by applying rotations of 75° and 125° about the corresponding axis, yielding 100 distinct initial attitude quaternions, to ensure comprehensive coverage of both moderate and large-angle slew maneuvers. The spacecraft and all reaction wheels are initialized with zero angular velocity, establishing rest-to-rest conditions with zero total angular momentum. All simulations incorporate representative environmental disturbances for low Earth orbit operations, including gravity-gradient torques and aerodynamic drag effects. Model mismatch is deliberately introduced through uncertainty in the spacecraft inertia tensor: the principal moments of inertia in the prediction model are perturbed with respect to the true values used in simulation, introducing structured parametric uncertainty in both diagonal and off-diagonal components. This structured model error creates prediction discrepancies that motivate residual learning compensation. Disturbance and uncertainty realizations remain fixed across paired simulations to isolate the effects of control architecture differences. Each simulation lasts 100 seconds, providing sufficient duration to observe both transient response and steady-state behavior. Convergence is declared when the attitude error magnitude remains below 0.5° for the remainder of the trajectory.

Controller Configurations: Three controllers are compared under identical conditions, including reference profiles, actuator constraints, and numerical integration settings, without controller retuning. The tested controllers are:

1. *Baseline NMPC*: Standard nonlinear model predictive control using the nominal prediction model without learning augmentation.
2. *NMPC + PINN*: NMPC augmented with PINN-based residual torque correction applied directly to the control input.
3. *NMPC + PINN + Supervisor*: PINN-augmented NMPC with residual corrections filtered through a Lyapunov-based supervisory layer prior to actuation.

© 2026 A. Ayanmanesh Motlaghmofrad et al., licenced via CC BY 4.0.

Metrics: Performance is quantified using two attitude-error-based metrics computed for each test realization:

- Steady-state RMSE: Root mean square attitude error computed over all time samples satisfying the 0.5° convergence criterion, characterizing steady-state precision.
- Final attitude error: Attitude error magnitude at the terminal simulation time, quantifying the final pointing accuracy.

For each controller, we aggregate metrics across all test runs, reporting mean values and standard deviations to characterize dispersion across initial conditions. Statistical significance is assessed through paired comparisons between each learning-augmented configuration and its corresponding baseline realization using the Wilcoxon signed-rank test. This nonparametric approach avoids normality assumptions and evaluates whether paired differences are systematically shifted from zero, thereby distinguishing consistent performance improvements from incidental gains on isolated trajectories. Additionally, we show the closed-loop time history of the *NMPC + PINN + Supervisor* controller to provide complementary qualitative insights into transient behavior and convergence characteristics.

4 Results

Aggregate statistical metrics computed over the full testing campaign are shown in Table 1. Mean steady-state RMSE values remain well below the prescribed 0.5° accuracy bound for all cases, indicating that observed variations occur within the high-precision regulation regime required by the task. Additionally, paired comparisons with respect to the baseline NMPC controller were performed using the Wilcoxon signed-rank test. Both learning-augmented configurations demonstrate statistically significant improvements relative to the baseline. The results support the hypothesis that PINN-based residual torque estimation enhances NMPC performance. The learned compensation systematically reduces both steady state RMSE and final attitude error across the test runs; a behavior that is consistent with the role of the PINN in correcting structured model mismatch, such as inertia uncertainty and unmodeled dynamic coupling, thereby causing the plant to more closely match the NMPC's internal predictive model. Moreover, improvements are not limited to typical trajectories: both the steady-state and final error dispersions decrease under learning-augmented control. This reduction in variance and improvement in terms of worst case (maximum error in Table 1), together with statistically significant paired improvements, demonstrates that the PINN enhances not only the nominal accuracy but also the uniformity and reliability of performance across diverse initial attitudes. While the supervised PINN configuration outperforms the baseline controller, its improvements are smaller than those of the unsupervised variant *NMPC + PINN*. Though, this outcome is expected, as the supervisory layer enforces a conservative Lyapunov-based non-increase condition using the nominal inertia model. As a result, beneficial corrections may be partially rejected when they appear unsafe under the nominal assumptions. The supervisor thus prioritizes stability over performance, yielding a controlled trade off: constrained use of the learned correction in exchange for preserving the guarantees of the baseline NMPC controller. Finally, Figure 2 shows the closed-loop attitude trajectories obtained under the hierarchical controller, which exhibit consistent convergence to the reference equilibrium.

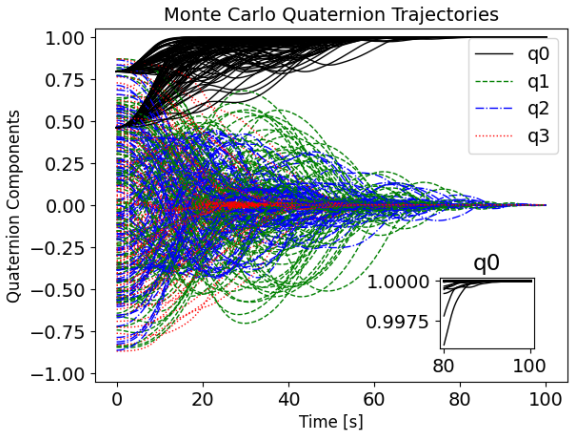


Figure 2: Test quaternion trajectories under NMPC with PINN-based residual compensation and Lyapunov-based supervisory filtering.

5 Discussion

This work presented a hierarchical control framework for attitude regulation of underactuated spacecraft that integrates NMPC, PINN-based residual compensation, and deterministic safety supervision. The nominal NMPC controller provides the primary control authority, handling nonlinear dynamics, actuator limits, and underactuation-aware maneuver planning. A physics-informed neural network is added as a residual torque estimator to compensate for model mismatch arising primarily from inertia uncertainty and neglected coupling effects. Finally, a Lyapunov-based supervisory layer ensures safety by filtering the learned correction through an energy-based admissibility condition. If the residual estimate is unreliable or rejected by the supervisor, the closed-loop system automatically reverts to the nominal NMPC policy without mode switching. This design aligns with safety-critical requirements typical of aerospace applications. The architecture was implemented in a high-fidelity simulation environment, accounting for real-time execution constraints, scheduling at multiple frequencies, numerical robustness, and consistency between training assumptions and closed-loop oper-

© 2026 A. Ayanmanesh Motlaghmofrad et al., licenced via CC BY 4.0.

Metric [deg]	*NMPC*	*NMPC PINN*	*NMPC PINN+Sup*
SS RMSE avg	0.1476 ± 0.0672	$\mathbf{0.1385 \pm 0.0437}$	0.1419 ± 0.0495
SS RMSE max	0.4940	$\mathbf{0.3392}$	0.3747
Final error avg	0.0870 ± 0.0704	$\mathbf{0.0722 \pm 0.0452}$	0.0759 ± 0.0497
Final error max	0.4915	$\mathbf{0.1607}$	0.2315

Table 1: Results obtained with 100 simulations using the three controller introduced in Section 3.

ating conditions. A test validation campaign showed that PINN-based residual compensation improves steady-state attitude accuracy and reduces performance dispersion. When used without supervision, the largest performance gains were observed; with the supervisory layer active, improvements remained statistically significant while preserving strict stability guarantees, highlighting a clear performance-safety trade-off. A relevant aspect concerns the computational and architectural complexity introduced by the proposed method relative to the observed performance gain. The PINN adds a feedforward evaluation at each NMPC sampling interval that does not affect the real-time feasibility of the $T_s = 0.5\,\mathrm{s}$ control loop, as shown in [12]. The Lyapunov supervisor further adds only a closed-form scalar admissibility check, incurring no additional optimization overhead. While the average improvement in steady-state RMSE is modest, the architecture yields larger gains in the worst-case and dispersion metrics (24.16% reduction in maximum steady-state RMSE, 29.42% reduction in final error variance). Given the low computational cost of the added components, we argue that even a modest average improvement is justified when accompanied by a substantial reduction in worst-case error.

6 Acknowledgments

This work has been developed with the contribution of the Politecnico di Torino Interdepartmental Centre for Service Robotics (PIC4SeR https://pic4ser.polito.it).

References

1. Crouch, P. Spacecraft attitude control and stabilization: Applications of geometric control theory to rigid body models. *IEEE Transactions on Automatic Control* **29,** 321–331 (1984).
2. Krishnan, H., McClamroch, N. H. & Reyhanoglu, M. Attitude stabilization of a rigid spacecraft using two momentum wheel actuators. *Journal of Guidance, Control, and Dynamics* **18,** 256–263 (1995).
3. Jin, L. & Li, Y. Model Predictive Control-Based Attitude Control of Under-Actuated Spacecraft Using Solar Radiation Pressure. *Aerospace* **9.** ISSN: 2226-4310 (2022).
4. Cena, C., Bucci, S. & Balossino, A. *Deep reinforcement learning for under-actuated satellite attitude control and reaction wheel desaturation using solar radiation pressure* in *Proceedings of the International Astronautical Congress (IAC)* (Oct. 2023).
5. Shirobokov, M., Trofimov, S. & Ovchinnikov, M. Survey of machine learning techniques in spacecraft control design. *Acta Astronautica* **186,** 87–97. ISSN: 0094-5765 (2021).
6. Celestini, D. *et al.* Transformer-Based Model Predictive Control: Trajectory Optimization via Sequence Modeling. *IEEE Robotics and Automation Letters* **9,** 9820–9827 (2024).
7. KrishnaKumar, K., Rickard, S. & Bartholomew, S. Adaptive neurocontrol for spacecraft attitude control. *Neurocomputing* **9,** 131–148. ISSN: 0925-2312 (1995).
8. Apolloni, B., Battini, F. & Lucisano, C. A co-operating neural approach for spacecrafts attitude control. *Neurocomputing* **16,** 279–307. ISSN: 0925-2312 (1997).
9. Biggs, J. D. & Fournier, H. Neural-Network-Based Optimal Attitude Control Using Four Impulsive Thrusters. *Journal of Guidance, Control, and Dynamics* **43,** 299–309 (2020).
10. Gu, W., Primatesta, S. & Rizzo, A. Physics-informed Neural Network for Quadrotor Dynamical Modeling. *Robotics and Autonomous Systems* **171,** 104569. ISSN: 0921-8890 (2024).
11. Falas, S., Asprou, M., Konstntinou, C. & Michael, M. K. Robust Power System State Estimation Using Physics-Informed Neural Networks. *IEEE Transactions on Industrial Informatics* **pre-print** (2025).
12. Cena, C., Martini, M. & Chiaberge, M. Learning robust satellite attitude dynamics with physics-informed normalizing flow. *Acta Astronautica* **245,** 970–981 (2026).
13. Varey, J., Ruprecht, J. D., Tierney, M. & Sullenberger, R. *Physics-Informed Neural Networks for Satellite State Estimation* in *2024 IEEE Aerospace Conference* (2024), 1–8.
14. Berberich, J. & Allgöwer, F. An Overview of Systems-Theoretic Guarantees in Data-Driven Model Predictive Control. *Annual Review of Control, Robotics, and Autonomous Systems* **8,** 77–100. ISSN: 2573-5144 (2025).
15. European Cooperation for Space Standardization (ECSS). *ECSS-E-HB-40-02A: Space engineering - Machine learning handbook. Handbook* 2024.

© 2026 A. Ayanmanesh Motlaghmofrad et al., licenced via CC BY 4.0.