Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Physics-aware adaptive gain scheduling for magnetic detumbling of a PocketQube satellite with embedded hardware validation

Muhammad Idriss Ishaq
Universiti Sains Malaysia
Norilmi Amilia Ismail
aenorilmi@usm.my
Universiti Sains Malaysia

Research Article

Keywords: Magnetic detumbling, B-dot control, PocketQube, Adaptive gain scheduling, Machine learning, Spacecraft ADCS

Posted Date: March 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8966174/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Physics-aware adaptive gain scheduling for magnetic detumbling of a PocketQube satellite with embedded hardware validation

Muhammad Ishaq Idriss · Norilmi Amilia Ismail
School of Aerospace Engineering, Universiti Sains Malaysia, 14300 Nibong Tebal, Pulau Pinang, Malaysia

ABSTRACT

Magnetic detumbling is a critical early-phase requirement for PocketQube satellites, and classical B-dot control remains widely used due to its simplicity and robustness. However, fixed-gain implementations are typically tuned conservatively to ensure stability across varying geomagnetic conditions, which can limit performance. This work introduces a physics-aware adaptive gain scheduling approach that augments the classical B-dot controller by modulating its scalar gain as a bounded function of geometry and energy-related state variables. A compact feedforward neural network is used to compute gain adjustments while preserving the original control structure.
The method is evaluated through a 200-case Monte Carlo campaign under realistic low Earth orbit magnetic-field conditions representative of spring-ejection deployment rates. Compared to a conservatively tuned fixed-gain baseline, the adaptive approach reduces mean detumbling time by 14.2%, with improvements most evident under unfavourable magnetic geometries. Real-time feasibility is demonstrated on flight-representative embedded hardware, confirming negligible computational overhead within a 10 Hz control cycle.
The results suggest that physics-aware gain adaptation can enhance detumbling performance while maintaining bounded behaviour and embedded-system compatibility, offering a practical route toward data-assisted control in resource-constrained small satellites.

Keywords: Magnetic detumbling · B-dot control · PocketQube · Adaptive gain scheduling · Machine learning · Spacecraft ADCS

Abbreviations

ADCS	Attitude Determination and Control System
ARM	Advanced RISC Machines processor architecture
BCET	Best-Case Execution Time
BF	Body-Fixed frame
CPU	Central Processing Unit
DSP	Digital Signal Processing
DWT	Data Watchpoint and Trace unit
ECI	Earth-Centred Inertial frame
ECEF	Earth-Centred, Earth-Fixed frame
EMA	Exponential Moving Average
FP32	32-bit Floating-Point precision
FPU	Floating-Point Unit
GPIO	General Purpose Input/Output
HIL	Hardware-in-the-Loop
IGRF	International Geomagnetic Reference Field
INT8	8-bit Integer precision
IQR	Interquartile Range
kB	Kilobyte
LEO	Low Earth Orbit
LQR	Linear Quadratic Regulator
MAE	Mean Absolute Error
MPC	Model Predictive Control
MSE	Mean Squared Error
NED	North–East–Down frame
NRLMSISE	Naval Research Laboratory Mass Spectrometer and

Incoherent Scatter Radar Exosphere model
RAM Random Access Memory
RK4 Fourth-Order Runge–Kutta integrator
RMSE Root Mean Square Error
SIMD Single Instruction, Multiple Data
SRP Solar Radiation Pressure
STM32 STMicroelectronics 32-bit microcontroller family
TFLite TensorFlow Lite
WCET Worst-Case Execution Time
WGS-84 World Geodetic System 1984
1P Single-unit PocketQube standard (50 × 50 × 50 mm³)

1 Introduction

1.1 Motivation and PocketQube ADCS context

PocketQube and other picosatellite platforms have emerged as a compelling paradigm for low-cost space missions, enabling rapid technology demonstration, distributed sensing, and broadened access to orbital environments. The PocketQube standard, defined by a 50 × 50 × 50 mm³ (1P) unit with a typical mass of approximately 250 g, represents the end of satellite miniaturisation, allowing complete spacecraft systems to be realised at mission costs well below those of conventional CubeSat platforms [1–4].

Recent PocketQube missions have demonstrated steadily increasing levels of attitude determination and control system (ADCS) sophistication. SMOG-P, launched in 2019, validated three-axis magnetic attitude control on a picosatellite platform (Budapest University, 2019). Delfi-PQ further demonstrated advanced subsystem integration and attitude-control concepts at the PocketQube scale [5], while the Unicorn-2 constellation showed the feasibility of coordinated multi-satellite operations for Earth observation missions under severe resource constraints.

European space agencies and research institutions have increasingly recognised PocketQube platforms as enabling technologies for distributed sensing, technology demonstration, and educational missions under constrained budgets [2, 6]. The operational success of these missions critically depends on rapid commissioning, for which autonomous detumbling capability represents a key subsystem requirement [7]. Extended detumbling durations increase the risk of power depletion during the early mission phase, potentially preventing payload activation and mission commissioning [8, 9].

Despite these advances, the extreme limitations imposed by PocketQube platforms in volume, mass, power, and onboard computation practically constrain ADCS design. The initial detumbling phase following deployment remains a critical operational bottleneck. Residual angular rates on the order of 20–60 deg/s are commonly reported for spring-loaded deployment mechanisms [10]. Such rates can saturate attitude sensors, misalign solar panels, and inhibit power generation. For 1P PocketQube, which typically provides only about 12.5 cm² of effective solar collection area, prolonged detumbling may lead to battery depletion within 48–72 hours [11]. Consequently, detumbling must be achieved rapidly and robustly under stringent computational constraints, with ADCS tasks often limited to a small fraction of available onboard processing resources [1].

For spacecraft operating under these constraints, magnetic actuation using air-core magnetorquers remains the only practical and widely adopted solution for detumbling. Magnetorquers require no propellant, prevent momentum saturation, and impose minimal mass penalties compared to reaction wheel systems [8]. The classical B-dot control law utilises measurements of the time derivative of the geomagnetic field to generate control torques that dissipate rotational kinetic energy without needing explicit angular velocity estimation [12]. Due to its simplicity, robustness, and very low computational requirement, B-dot control has extensive flight heritage and continues to serve as the baseline detumbling strategy for CubeSat- and PocketQube-class missions [1, 13]. Consequently, detumbling must be achieved rapidly and robustly under stringent computational constraints, with ADCS tasks often limited to a small fraction of available onboard processing resources [14].

1.2 Classical B-dot control and geometric limitations

The B-dot control law, originally introduced by [15] and later rigorously analysed by [16, 17], commands a magnetic dipole moment proportional to the time derivative of the measured geomagnetic field, $\mathrm{m} = -k\,\dot{\mathrm{B}}$ where $k > 0$ is a scalar control gain. The resulting magnetic torque $\tau = \mathrm{m} \times \mathbf{B}$ dissipates rotational kinetic energy at a rate proportional to $\| \omega \times \mathbf{B} \|^2$, yielding Lyapunov stability under the assumption of a slowly varying geomagnetic field [16].

Despite its proven robustness and extensive flight heritage, fixed-gain B-dot control exhibits well-documented limitations arising from its intrinsic geometric dependence. Magnetic control authority is maximised when the angular velocity vector $\boldsymbol{\omega}$ is perpendicular to the geomagnetic field B, but vanishes entirely when the two vectors are aligned due to the cross-product structure of magnetic torque [18]. As a result, detumbling convergence rates may vary by an order of magnitude depending on initial attitude conditions, orbital inclination, and magnetic field geometry [17].

To guarantee stability under worst-case near-parallel alignment, fixed-gain B-dot controllers must be tuned conservatively, which significantly degrades performance under favourable geometric conditions. Several studies report that such conservative tuning sacrifices approximately 30–50% of achievable detumbling performance compared to geometry-aware gain selection [19]. This fundamental trade-off between guaranteed stability and attainable performance has motivated sustained research into adaptive and gain-scheduled magnetic detumbling strategies.

1.2 Related work and research gap

Alternative magnetic control laws, such as Y-dot [20] and M-dot, have been proposed, each with specific advantages for different mission scenarios, though B-dot remains the most widely adopted for detumbling due to its simplicity and robustness. A variety of adaptive and time-varying magnetic detumbling approaches have been proposed to mitigate geometric performance degradation while preserving stability guarantees. Neural adaptive control has been applied to spacecraft attitude problems [21], though most approaches target pointing tasks with reaction wheels rather than magnetic detumbling. The geometry-aware gain scaling strategy introduced by [12] demonstrated improved detumbling performance in simulation, while Wailand and Bauer (2020) [22] examined the influence of gain selection and sensor noise on magnetic detumbling behaviour. Recent work on magnetic attitude stabilisation with reduced-actuator configurations [23] has demonstrated that stable control can be achieved even with under-actuated systems. However, detumbling remains more challenging due to unknown initial conditions. However, many such approaches rely on heuristic tuning rules, offline optimisation, or simplified simulation assumptions, and typically lack validation on flight-representative embedded hardware.

More recently, advances in machine learning have motivated increasing interest in neural-network-based spacecraft attitude control and adaptation. Machine learning approaches to spacecraft control have gained significant attention, with methods broadly categorised into supervised learning and reinforcement learning [24]. Reinforcement learning approaches have demonstrated substantial performance improvements in simulation environments [25, 26], while deep learning techniques have been applied to attitude estimation and onboard fault detection [27]. Nevertheless, these methods face significant barriers to deployment on PocketQube-class spacecraft. In particular, most learning-based controllers do not preserve the Lyapunov stability structure of classical magnetic control laws, require computational resources exceeding those available on flight-qualified microcontrollers, and are validated exclusively in software simulation without accounting for quantisation effects, timing jitter, or real-time execution constraints [28, 29].

Recent work has demonstrated that compact neural networks can replicate classical adaptive B-dot control laws on microcontroller-class hardware with minimal computational cost. However, controller replication alone does not overcome the fundamental performance limitations imposed by conservative gain tuning. The present study advances this line of research by formulating a physics-aware adaptive gain scheduling framework that leverages learned geometric and energetic dependencies to improve detumbling performance, while explicitly preserving the energy dissipation guarantees of classical magnetic control. A structured comparison of classical, adaptive, and learning-based magnetic detumbling approaches is provided in Table 1, highlighting the lack of physics-aware learning with flight-class embedded validation.

Table 1 Comparison Of Magnetic Detumbling Control Approaches

Reference	Method	Adaptivity	Physics-Based	Learning	Hardware Validation	Key Limitation
Stickler & Alfriend (1976) [15]	Classical B-dot	✗	✓	✗	✗	Very conservative gains
Lovera & Astolfi (2004) [16]	Energy-based magnetic control	✗	✓✓	✗	✗	Averaged dynamics assumption
Psiaki (2001) [30]	Periodic LQR magnetic control	Implicit	✓✓	✗	✗	Offline design, no embedded validation
Desouky & Abdelkhalik (2020) [12]	Modified B-dot	Partial	✓	✗	✗	Simulation-only tuning
Petit & Sarras (2020) [31]	Lyapunov-based law	✗	✓✓	✗	✗	Conservative gain selection
Monkell et al. (2018) [32]	Reduced-actuator B-dot	✗	✓	✗	✗	Slower convergence
Farissi et al. (2019) [13]	B-dot + HIL	✗	✓	✗	✓	Fixed-gain control
Ahmed & Kerrigan (2014) [33]	MPC detumbling	✓	✓	✗	✗	Computational burden
This work	Physics-aware adaptive B-dot	✓ (bounded)	✓✓	✓ (supervised, heuristic labels)	✓ (STM32, embedded validation)	Empirical stability only; heuristic-based training

1.3 Contributions and paper organisation

This paper presents a flight-ready adaptive gain scheduling subsystem for magnetic detumbling, designed for integration into PocketQube ADCS architectures. The work demonstrates that bounded data-driven adaptation can enhance heritage magnetic control performance while satisfying the severe resource constraints of picosatellite avionics. The principal contributions to space systems engineering practice are as follows:

1. A physics-aware adaptive gain scheduling framework that modulates the scalar gain of the classical B-dot controller using 13 physically meaningful features, including angular velocity–magnetic field geometry, magnetic field strength, and energy dissipation indicators. The framework maintains bounded gain adaptation to preserve empirically verified energy dissipation behaviour.
2. A high-fidelity Monte Carlo evaluation dataset comprising 200 randomised detumbling scenarios under realistic orbital and geomagnetic conditions, yielding 2.88 million validated training samples derived from first-principles rigid-body dynamics and magnetic actuation models.
3. Comprehensive embedded implementation and validation on a flight-class STM32 microcontroller, including cycle-accurate timing analysis, memory and power profiling, and embedded hardware testing demonstrating real-time feasibility under PocketQube resource constraints.
4. Statistical and geometry-dependent performance analysis showing improved convergence rates and revealing that adaptive gains yield the greatest benefit under unfavourable magnetic geometries, where fixed-gain B-dot control is inherently conservative.
5. An open-source, reproducible implementation encompassing simulation, training, and embedded deployment artefacts to support independent validation and future extension.

The remainder of this work is organised as follows. Section 2 establishes the theoretical foundations for spacecraft dynamics, geomagnetic modelling, and classical B-dot control. Section 3 introduces the proposed adaptive gain scheduling framework. Section 4 describes the training methodology and dataset generation. Section 5 details the embedded implementation and embedded hardware validation. Section 6 presents simulation results and statistical analyses. Section 7 discusses limitations and implications for flight deployment, and Section 8 concludes the work.

2 Dynamical model

This section establishes the mathematical foundation for simulating the attitude dynamics of PocketQube, orbital motion, environmental disturbances, and magnetic control actuation. The modelling framework incorporates rigid-body dynamics, high-fidelity geomagnetic field representation, and realistic actuator constraints to provide a physically consistent testbed for evaluating the proposed physics-aware adaptive gain scheduling approach.

2.1 Reference frames

To model spacecraft motion, environmental interactions, and control behaviour consistently, the following reference frames are employed [34–36]:

Earth-Centred Inertial (ECI) Frame: A non-rotating inertial reference frame centred at the Earth's centre of mass. The ECI frame is used for orbital propagation, attitude dynamics, and modelling inertial forces. Its axes are defined with the X-axis pointing toward the vernal equinox, the Z-axis aligned with Earth's rotation axis, and the Y-axis completing a right-handed triad.

Body-Fixed (BF) Frame: A rotating frame rigidly attached to the spacecraft, with its axes aligned with the spacecraft's principal moments of inertia. All onboard sensors (magnetometers, gyroscopes), actuators (magnetorquers), control laws (B-dot and adaptive gain scheduling), and torque computations are expressed in this frame. The origin coincides with the spacecraft's centre of mass.

Earth-Centred, Earth-Fixed (ECEF) Frame: A rotating frame fixed to Earth's surface, rotating with angular velocity $\omega_{\oplus} = 7.2921 \times 10^{-5}$ rad/s. This frame serves as an intermediate coordinate system for converting inertial positions to geodetic coordinates required by the geomagnetic field model.

Geodetic Frame (Latitude–Longitude–Altitude): Defines the spacecraft position relative to the WGS-84 ellipsoid using latitude ϕ, longitude λ, and altitude h. This representation is required for evaluating the International Geomagnetic Reference Field (IGRF).

North–East–Down (NED) Frame: A local geographic frame aligned with the Earth's surface at the spacecraft position. The IGRF model outputs the magnetic field vector in this frame, which is subsequently transformed to ECI and then to the body frame for control computation. The reference frames adopted for attitude dynamics, geomagnetic field evaluation, and magnetic control are illustrated in Fig. 1.

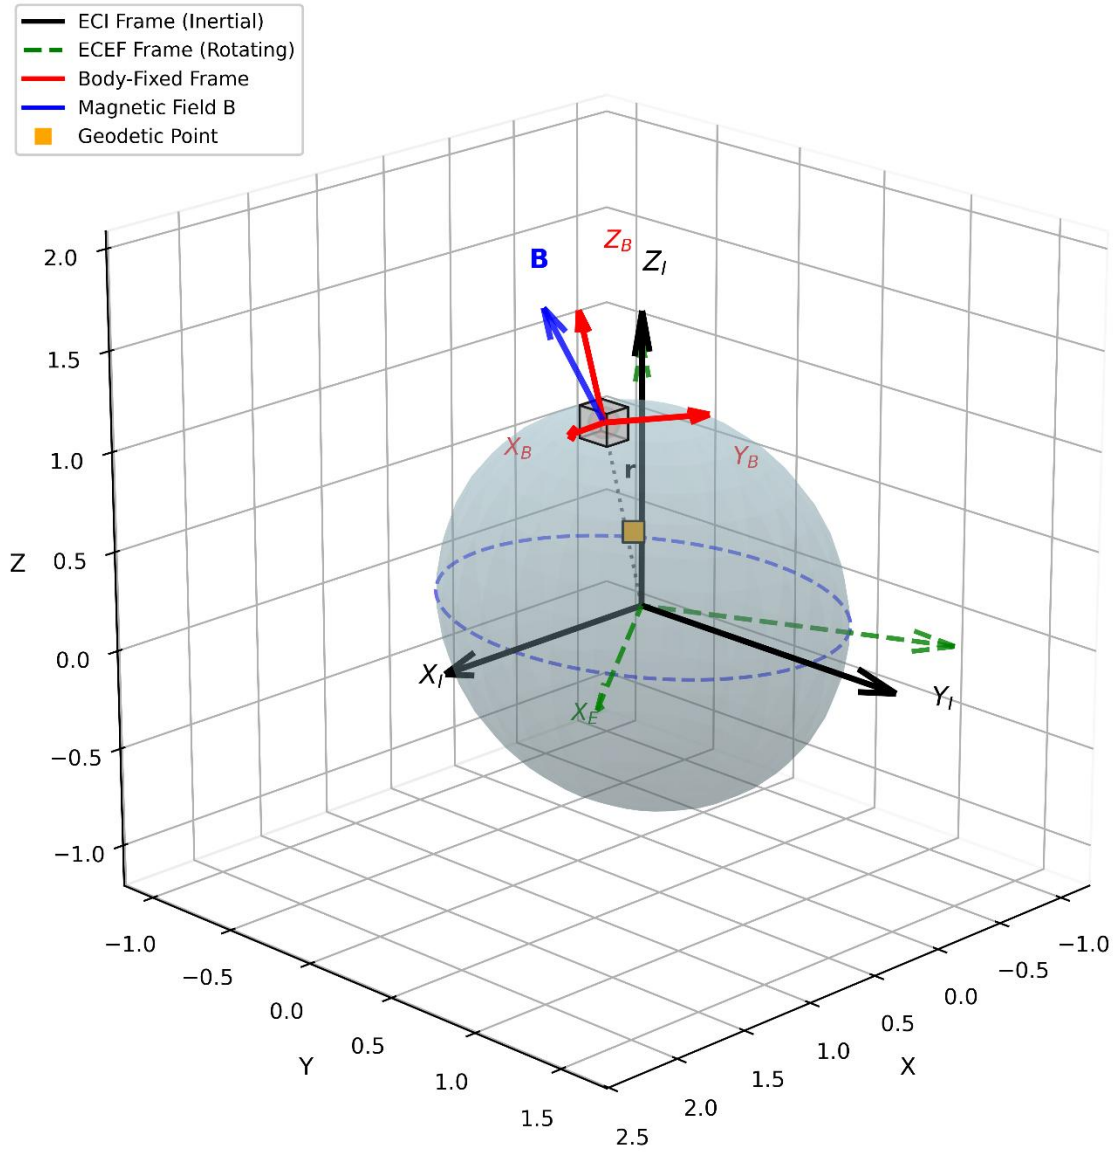


Fig. 1. Definition of the reference frames used.

2.2 Spacecraft rigid-body dynamics

2.2.1 Physical configuration

The spacecraft is modelled as a rigid 1P PocketQube with mass $m = 0.25$ kg and dimensions $L \times W \times H = 0.05 \times 0.05 \times 0.05$ m.

Assuming uniform density and cubic geometry, the inertia tensor about the principal axes is diagonal and given by

$$\mathbf{I} = \mathrm{diag}(I_x, I_y, I_z) \tag{2}$$

with

$$I_x = I_y = I_z = 0.00125\ \mathrm{kg{\cdot}m^2} \tag{3}$$

The symmetric inertia simplifies rotational dynamics while remaining representative of typical PocketQube mass distributions.

The reference PocketQube spacecraft configuration adopted for this study is summarised in Table 2, representative of 1P-class platforms with magnetic attitude control capability.

Table 2 Reference PocketQube Spacecraft Parameters

Parameter	Value	Units
Orbit		
Altitude	550	km
Inclination	97.4	deg
Eccentricity	0 (circular)	—
Orbital period	95.8	min
Spacecraft		
Configuration	1P PocketQube	—
Mass	0.25	kg
Dimensions	5 × 5 × 5	cm
Inertia tensor	diag(0.00125, 0.00125, 0.00125)	kg·m²
Magnetorquers		
Maximum dipole	±0.1 (all axes)	A·m²
Hysteresis lag	1.0	s
Switching threshold	1.0×10^{-4}	A·m²
Control		
Update rate	10	Hz
Baseline gain	10.0	A·m²·s/T
EMA filter coefficient	0.18	—
Angular velocity smoothing	0.85	—
Monte Carlo		
Number of scenarios	200	—
Initial angular velocity	±30 (random)	deg/s
Convergence threshold	0.005	rad/s
Maximum simulation time	10,000	s
Environment		
Magnetic field model	IGRF-13 (epoch 2024.5)	—
Atmospheric density (550 km)	2.0×10^{-13}	kg/m³
Solar radiation pressure	4.56×10^{-6}	N/m²
Earth's angular velocity	7.2921×10^{-5}	rad/s
Numerical Integration		
Integration method	RK4 (fixed-step)	—
Timestep	0.1	s
Quaternion normalization	Every step	—

2.2.2 Euler rotational dynamics

The rotational motion of the spacecraft is governed by Euler's equation [34–36]:

$$\mathbf{I}\dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (\mathbf{I}\boldsymbol{\omega}) = \boldsymbol{\tau} \quad (4)$$

where $\boldsymbol{\omega}$is the angular velocity expressed in the body frame and $\boldsymbol{\tau}$is the total external torque. Solving for angular acceleration yields

$$\dot{\boldsymbol{\omega}} = \mathbf{I}^{-1}(\boldsymbol{\tau} - \boldsymbol{\omega} \times (\mathbf{I}\boldsymbol{\omega})) \quad (5)$$

For a diagonal inertia tensor aligned with the principal axes, Euler's equations decouple into three scalar equations. For the symmetric 1P PocketQube configuration, the gyroscopic coupling terms vanish, reducing computational complexity without compromising physical fidelity.

2.2.3 Quaternion attitude kinematics

Spacecraft attitude is represented using a unit quaternion $\mathbf{q} = [q_0, \mathbf{q}_v]^T$, where $q_0$is the scalar component and $\mathbf{q}_v$is the vector component. The quaternion kinematic equation is given by [34, 35]:

$$\dot{\mathbf{q}} = \frac{1}{2}\boldsymbol{\Omega}(\boldsymbol{\omega})\mathbf{q} \quad (6)$$

where $\boldsymbol{\Omega}(\boldsymbol{\omega})$ is the skew-symmetric quaternion rate matrix. Quaternion normalisation is enforced at each timestep to prevent numerical drift:

$$\mathbf{q} \leftarrow \frac{\mathbf{q}}{\|\mathbf{q}\|} \quad (7)$$

This representation avoids singularities associated with Euler angles and is well-suited for real-time embedded implementation.

2.3 Orbital dynamics and geomagnetic field modelling

2.3.1 Orbital propagation

The spacecraft operates in a circular sun-synchronous orbit with parameters summarised in Table 2. The orbital radius is computed as [35]:

$$r = R_{\oplus} + h \quad (8)$$

and the orbital angular velocity is:

$$\omega_o = \sqrt{\frac{\mu}{r^3}} \quad (9)$$

where μ is Earth's gravitational parameter. The spacecraft's position in the ECI frame is propagated using classical two-body dynamics.

2.3.2 International Geomagnetic Reference Field (IGRF-13)

The Earth's magnetic field is modelled using the IGRF-13 model, valid for epochs 2020–2025 [37]. The model represents the geomagnetic field as a spherical harmonic expansion up to degree 13, providing accuracy suitable for attitude control studies.

At each timestep, the geomagnetic field is evaluated using the spacecraft's geodetic coordinates (ϕ, λ, h) and returned in the NED frame. The field vector is subsequently transformed through the sequence: ECI to ECEF to Geodetic to NED to ECEF to ECI to Body frame.

For the specified orbit, the magnetic field magnitude varies between approximately 25 μT and 52 μT, with secular variation negligible relative to orbital timescale variation.

2.4 Environmental disturbance torques

Three environmental disturbance torques are included: gravity-gradient, solar radiation pressure (SRP), and aerodynamic drag. Residual magnetic dipoles from onboard electronics can also contribute to disturbance torques in miniature satellites [38], though these are neglected in the present model. All torques are computed in the body frame and summed with the control torque.

2.4.1 Gravity-gradient torque

The gravity-gradient torque is given by [34–36]:

$$\boldsymbol{\tau}_{gg} = 3\frac{\mu}{r^3}\hat{\mathbf{r}} \times \mathbf{I}\hat{\mathbf{r}} \tag{10}$$

For the symmetric inertia tensor considered here, gravity-gradient torque is empirically zero. Residual asymmetries may introduce torques on the order of 10^{-9} Nm, which are negligible during detumbling.

2.4.2 Solar radiation pressure torque

Solar radiation pressure torque is modelled as [34–36]:

$$\boldsymbol{\tau}_{srp} = P_{\odot} C_r A(\mathbf{r}_{cp} \times \hat{\mathbf{s}}) \tag{11}$$

where standard values are used for solar pressure, reflectivity, and centre-of-pressure offset. Eclipse conditions are explicitly modelled, disabling SRP when the spacecraft is in Earth's shadow.

2.4.3 Aerodynamic drag torque

Aerodynamic drag torque is modelled using atmospheric density from the NRLMSISE-00 model [34–36]. At 550 km altitude, drag torque magnitudes are on the order of 10^{-10} Nm, remaining negligible compared to the magnetic control torque.

2.4.4 Total external torque

The net torque applied to the spacecraft is:

$$\boldsymbol{\tau} = \boldsymbol{\tau}_c + \boldsymbol{\tau}_{gg} + \boldsymbol{\tau}_{srp} + \boldsymbol{\tau}_{aero} \tag{12}$$

During detumbling, magnetic control torque dominates, with disturbances contributing less than 0.15% of the total.

2.5 Magnetorquer actuator model

Three orthogonal air-core magnetorquers generate a magnetic dipole moment $\mathbf{m}$, producing control torque [35]:

$$\boldsymbol{\tau}_c = \mathbf{m} \times \mathbf{B} \tag{13}$$

Dipole moments are constrained by actuator saturation and clipped to hardware limits. A hysteresis switching model is implemented to prevent rapid sign reversals and actuator chattering [39]. Electrical power consumption is computed from coil resistance and commanded current, yielding a peak power of approximately 2 mW per axis.

2.6 Classical B-dot control law

The baseline detumbling controller employs the classical B-dot algorithm [12, 37] [38,39]. The time derivative of the magnetic field is approximated using finite differencing and filtered using an exponential moving average to suppress sensor noise.

To ensure control torque remains orthogonal to the geomagnetic field, the derivative is projected onto the plane normal to $\mathbf{B}$. For very small angular rates, angular velocity smoothing is applied to prevent numerical instability. The commanded magnetic dipole moment is

$$\mathbf{m} = -k_{\text{base}}\dot{\mathbf{B}}_{\perp} \tag{14}$$

subject to saturation and hysteresis constraints. The negative sign ensures energy dissipation by opposing rotational motion.

2.6.1 Baseline gain selection

For the specified 550 km sun-synchronous orbit, the orbital-average geomagnetic field magnitude predicted by the IGRF-13 model is approximately 28 μT, consistent with typical low-Earth orbit field strengths [40, 41]. This value is used solely as a normalisation reference for baseline gain selection and does not enter the stability analysis.

The baseline B-dot gain is selected following standard magnetic detumbling practice, accounting for the discrete control update period (Δt = 0.1 s) and a conservative safety factor to prevent actuator saturation and preserve negative energy dissipation under unfavourable magnetic geometries [16]. The resulting fixed gain is

$$k_{\text{base}} = 10.0\ \text{Am}^2\text{s/T} \tag{15}$$

Such conservative gain selection is widely adopted to guarantee stability across all magnetic field orientations, at the expense of reduced performance in favourable geometries [18].

This conservative selection maintains bounded response across all magnetic field geometries at the expense of reduced performance under favourable conditions, motivating adaptive gain scheduling.

2.7 Simulation parameters

Table 3 summarises the simulation parameters used for evaluating the proposed adaptive gain scheduling framework. The complete set of thirteen physics-aware input features is summarised in Table 3.

Table 3 Physics-Aware Feature Summary and Interpretations

N	Symbol	Description	Physical Interpretation	Typical Range
1–3	$\omega_x, \omega_y, \omega_z$	Angular velocity components	Rotational state, energy proxy	±0.5 rad/s
4–6	B_x, B_y, B_z	Magnetic field components	Control direction, orbital position	±52 μT
7	$\| \boldsymbol{\omega} \|$	Angular rate magnitude	Kinetic energy: $E = \frac{1}{2} I \| \omega \|^2$	0–0.6 rad/s
8	$\| \mathbf{B} \|$	Magnetic field magnitude	Control authority scaling	25–52 μT
9	$\cos(\boldsymbol{\omega}, \mathbf{B})$	Geometry alignment	Control efficiency indicator	-1 to +1
10	dE/dt	Energy dissipation rate	Lyapunov stability metric	-0.01 to 0 W
11	$\tau_{\max}$	Maximum torque available	Saturation awareness	2.5–5.2 × 10^{-6} N·m
12	h	Orbital altitude	Environmental context	549–551 km
13	$\| d\mathbf{B}/dt \|$	Field derivative magnitude	Observability metric	0–1 × 10^{-4} T/s

2.8 Validation and conservation checks

To ensure numerical fidelity and physical consistency, angular momentum drift, quaternion normalisation error, and rotational energy dissipation are monitored at every integration step. Simulation runs violating predefined conservation or normalisation thresholds are excluded from subsequent statistical analysis.

Formal Lyapunov stability proofs for magnetic detumbling typically assume constant control gains. The present study does not provide analytical stability proofs for the time-varying gain case. Instead, bounded energy dissipation under adaptive gain modulation is verified empirically through extensive time-domain simulation across 200 randomised scenarios.

3. Physics-aware adaptive gain scheduling framework

This section presents the proposed physics-aware adaptive gain scheduling framework, which constitutes the central contribution of this work. Rather than replacing the classical B-dot magnetic detumbling law, the proposed method augments it with a learned, bounded gain-modulation mechanism that responds to instantaneous geometric and energetic conditions while preserving the energy-dissipation properties established in Section 2.

The framework is designed to retain the interpretability, robustness, and flight heritage of classical magnetic attitude control while selectively exploiting data-driven learning to improve performance under favourable operating conditions. In contrast to end-to-end learning approaches that map raw sensor measurements directly to actuator commands, the proposed architecture embeds physical structure, stability constraints, and actuator limitations explicitly into the control loop.

3.1 Design philosophy and control architecture

Recent learning-based spacecraft control studies increasingly employ neural networks as black-box function approximators trained end-to-end to generate control commands from sensor data [25, 26]. While such approaches can demonstrate strong simulation performance, they often lack interpretability, formal safety guarantees, and compatibility with flight-qualified control architectures, properties that remain essential for operational spacecraft systems. The overall control architecture is shown in Fig. 2.

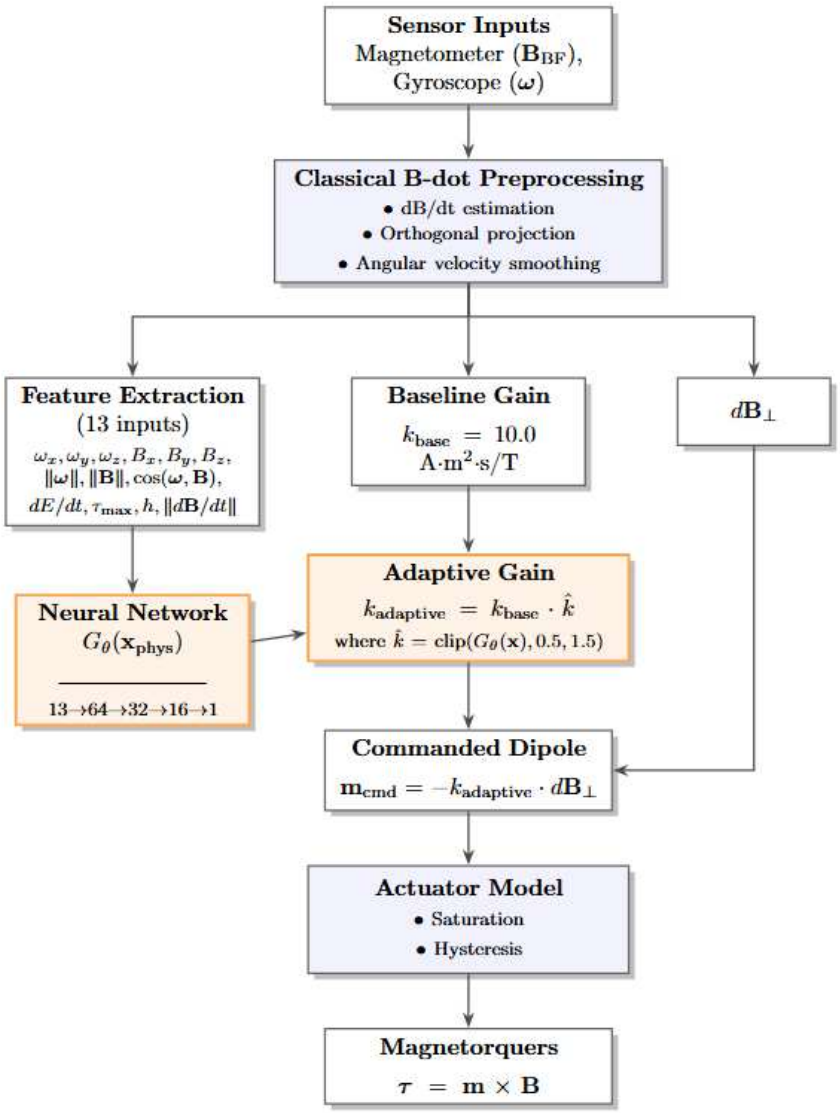


Fig. 2. Block diagram of the physics-aware adaptive magnetic detumbling controller.

In contrast, this work adopts a *physics-aware learning paradigm*, in which classical control structure is preserved, and learning is confined to a strictly bounded adaptation mechanism. Specifically, the adaptive component modulates only the scalar gain of the classical B-dot control law, maintaining the original control direction and energy dissipation behaviour. The commanded magnetic dipole moment is given by

$$\mathbf{m}_{\mathrm{cmd}}(t) = -k_{\mathrm{adaptive}}(t)\,\dot{\mathbf{B}}_{\perp}(t) \tag{16}$$

where $\dot{\mathbf{B}}_{\perp}$denotes the magnetic field derivative projected onto the plane orthogonal to the geomagnetic field (Eqns. 37–40), and $k_{\mathrm{adaptive}}(t)$is a time-varying adaptive gain.

The adaptive gain is parameterised as a modulation of a fixed, classically tuned baseline gain,

$$k_{\mathrm{adaptive}}(t) = k_{\mathrm{base}}\,\hat{k}(t) \tag{17}$$

where $k_{\mathrm{base}} = 10.0\ \mathrm{A \cdot m^2 \cdot s/T}$(Eqn. 58), and $\hat{k}(t)$is a dimensionless modulation factor learned by a neural network and constrained within predefined bounds.

This formulation provides several important properties:

- Baseline stability preservation: Setting $\hat{k}(t) \equiv 1$ recovers the classical fixed-gain B-dot controller.
- Bounded adaptation: Enforcing $\hat{k}(t) \in [0.5,1.5]$prevents destabilising gain excursions and actuator saturation.
- Interpretability: The adaptive term has a clear physical meaning as a multiplicative gain correction, facilitating operator insight and fault diagnosis.

The adaptive modulation factor is computed as

$$\hat{k}(t) = G_{\theta}(\mathbf{x}_{\mathrm{phys}}(t)) \tag{18}$$

where $G_{\theta}: \mathbb{R}^{13} \to \mathbb{R}$is a feedforward neural network with parameters θ, and $\mathbf{x}_{\mathrm{phys}}$is a vector of physics-aware features described in Section 3.2.

3.2 Physics-aware feature construction

The effectiveness and reliability of the proposed framework depend critically on the choice of neural network inputs, rather than providing raw sensor measurements and relying on the network to infer physical relationships implicitly. This work explicitly encodes domain knowledge from spacecraft dynamics, magnetic actuation theory, and geometric control into a structured feature set. The physics-aware feature vector is defined as:

$$\mathbf{x}_{\mathrm{phys}}(t) = \begin{bmatrix} \boldsymbol{\omega}_{\mathrm{sm}}(t) \\ \mathbf{B}_{\mathrm{BF}}(t) \\ \mathbf{f}_{\mathrm{derived}}(t) \end{bmatrix} \in \mathbb{R}^{13} \tag{19}$$

where $\boldsymbol{\omega}_{\mathrm{sm}} \in \mathbb{R}^3$denotes smoothed angular velocity components, $\mathbf{B}_{\mathrm{BF}} \in \mathbb{R}^3$is the body-frame geomagnetic field vector, and $\mathbf{f}_{\mathrm{derived}} \in \mathbb{R}^7$contains physics-derived scalar features.

The gain scheduler is realised using a fully connected feedforward neural network with architecture specified in Fig. 3, comprising 3,521 trainable parameters suitable for deployment on microcontroller-class hardware. The gain scheduler is implemented using the compact neural network shown in Fig. 3.

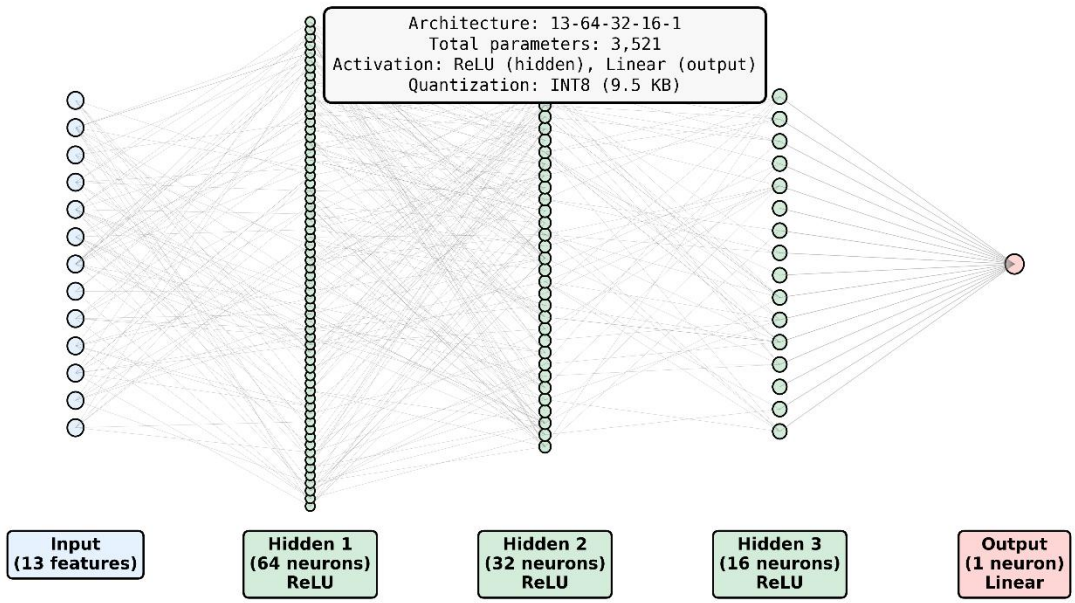


Fig. 3. Neural network architecture used for adaptive gain scheduling.

3.2.1 Angular velocity and magnetic field components

Components 1–3: Angular velocity components

The smoothed angular velocity vector $\boldsymbol{\omega}_{\text{sm}} = [\omega_x, \omega_y, \omega_z]^T$(rad/s) provides direct information about the spacecraft's rotational state. Smoothing is applied using Eqn. 41 to reduce noise amplification near convergence when $\| \boldsymbol{\omega} \| < 0.005$ rad/s. These components allow the network to adapt gains based on both the magnitude and distribution of angular momentum.

Components 4–6: Magnetic field components

The body-frame magnetic field vector $\mathbf{B}_{\text{BF}} = [B_x, B_y, B_z]^T$(T) conveys information about both orbital position and instantaneous control authority. Since the magnetic control torque is constrained to lie perpendicular to **B** The orientation of this vector relative to the body frame directly influences the achievable damping directions [15, 30].

3.2.2 Energy and control authority scalars

Component 7: Angular velocity magnitude

$$\| \boldsymbol{\omega} \| = \sqrt{\omega_x^2 + \omega_y^2 + \omega_z^2} \tag{20}$$

which serves as a proxy for rotational kinetic energy. For the symmetric PocketQube considered here, $E = \frac{1}{2} I_0 \| \boldsymbol{\omega} \|^2$, enabling energy-aware gain modulation [36].

Component 8: Magnetic field magnitude

$$\| \mathbf{B} \| = \sqrt{B_x^2 + B_y^2 + B_z^2} \tag{21}$$

Along the selected sun-synchronous orbit, $\| \mathbf{B} \|$varies by approximately a factor of two between equatorial and polar regions [37]. Because achievable magnetic torque scales linearly with field strength, this feature enables adaptive exploitation of favourable geomagnetic conditions [17, 42].

3.2.3 Geometric control efficiency indicator

Component 9: Alignment between angular velocity and magnetic field

$$\cos(\boldsymbol{\omega}, \mathbf{B}) = \frac{\boldsymbol{\omega} \cdot \mathbf{B}}{\|\boldsymbol{\omega}\| \, \|\mathbf{B}\|} \tag{22}$$

This scalar captures the fundamental geometric limitation of magnetic control. When $\boldsymbol{\omega}$is perpendicular to $\mathbf{B}$ Control authority is maximised; when they are aligned, the achievable torque approaches zero. Classical fixed-gain B-dot controllers must be tuned conservatively for worst-case alignment [17, 19, 43]. Providing this feature explicitly allows the network to learn geometry-aware gain modulation.
The instantaneous energy dissipation rate satisfies

$$P(t) = -k \parallel \boldsymbol{\omega} \times \mathbf{B} \parallel^2 = -k \parallel \boldsymbol{\omega} \parallel^2 \parallel \mathbf{B} \parallel^2 (1 - \cos^2(\boldsymbol{\omega}, \mathbf{B})) \quad (23)$$

highlighting the direct influence of geometry on detumbling efficiency.

3.2.4 Energy dissipation and saturation awareness

Component 10: Instantaneous energy dissipation rate

$$\frac{dE}{dt} = -k_{\text{base}} \parallel \boldsymbol{\omega}_{\text{sm}} \times \mathbf{B}_{\text{BF}} \parallel^2 \quad (24)$$

This quantity encodes Lyapunov-based stability information and provides a physics-consistent indicator of control effectiveness [43].
Component 11: Maximum available magnetic torque

$$\tau_{\max}(t) = m_{\max} \parallel \mathbf{B}_{\text{BF}}(t) \parallel \quad (25)$$

where $m_{\max} = 0.1\ \text{A} \cdot \text{m}^2$. This feature enables saturation-aware gain modulation and avoids energy waste when commanded torques are unachievable [12, 44–46].

3.2.5 Environmental and observability context

Component 12: Orbital altitude

$$h = \parallel \mathbf{r}_{\text{ECI}} \parallel - R_{\oplus} \quad (26)$$

Although nearly constant for the target orbit, altitude provides context for magnetic field modelling and residual environmental torques [47, 48].
Component 13: Magnetic field derivative magnitude $\parallel \frac{d\mathbf{B}_{\text{BF}}}{dt} \parallel$.

This feature serves as an observability indicator, reflecting how well angular velocity is captured through magnetometer measurements under the slowly varying field assumption [7].

3.3 Neural network architecture

The adaptive gain function $G_\theta$is implemented as a compact feedforward neural network explicitly designed for embedded deployment. The architecture balances representational capacity, computational efficiency, and numerical robustness under INT8 quantisation.
The network consists of four fully connected layers with dimensions $13 \to 64 \to 32 \to 16 \to 1$, using hyperbolic tangent activations in the hidden layers:

$$\mathbf{h}^{(1)} = \tanh(\mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)}) \quad (27)$$
$$\mathbf{h}^{(2)} = \tanh(\mathbf{W}^{(2)}\mathbf{h}^{(1)} + \mathbf{b}^{(2)})$$
$$\mathbf{h}^{(3)} = \tanh(\mathbf{W}^{(3)}\mathbf{h}^{(2)} + \mathbf{b}^{(3)})$$
$$\hat{k}_{\text{raw}} = \mathbf{w}^{(4)}\mathbf{h}^{(3)} + b^{(4)} \quad (28)$$

The output is clipped to enforce bounded adaptation,

$$\hat{k}(t) = \text{clip}(\hat{k}_{\text{raw}}, 0.5, 1.5) \quad (29)$$

The total number of trainable parameters is 3,521, corresponding to 14.1 kB in FP32 and approximately 9.5 kB after INT8 quantisation [49].

3.4 Bounded adaptation and empirical stability verification

Introducing time-varying gains into magnetic detumbling and certain under-actuated configurations [23] complicates classical Lyapunov analysis [17, 50]. Formal stability proofs for sampled-data, quantised, time-varying gains remain an open theoretical problem [51]. While we do not provide a formal Lyapunov proof, this work enforces design constraints intended to prevent destabilisation:

1. Strictly positive bounded gains: $5.0 \leq k_{\text{adaptive}}(t) \leq 15.0$.
2. Smooth temporal variation: The 10 Hz update rate and smooth activations limit the gain rate-of-change.
3. Safe fallback capability: Setting $\hat{k} \equiv 1$ recovers the proven fixed-gain controller.

Empirical stability is verified through:

- Negative energy dissipation (dE/dt < 0) across all Monte Carlo scenarios
- Bounded angular velocity trajectories in all test cases
- No divergent behaviour observed during embedded validation.

These empirical checks provide confidence in operational stability, though formal guarantees remain a topic for future theoretical work.

3.5 Training objective

Network parameters are optimised using a supervised regression objective,

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{MSE}} + \lambda \parallel \theta \parallel_2^2 \tag{30}$$

where target gains are generated using physics-motivated pseudo-labels derived from baseline energy dissipation trends rather than analytical optimal control solutions [52].

4. Training methodology and dataset generation

This section describes the training methodology for the adaptive gain scheduler, including Monte Carlo dataset generation, physics-motivated label construction, and neural network training. The approach prioritises physical consistency, prevents data leakage through scenario-level partitioning, and integrates embedded deployment constraints from the outset.

4.1 Monte Carlo simulation framework

Training data were generated through high-fidelity Monte Carlo simulations of PocketQube detumbling dynamics using the classical fixed-gain B-dot controller described in Section 2 as a reference baseline. Unlike learning approaches trained on simplified or surrogate models, the same simulation framework employed for final performance evaluation was used for dataset generation, ensuring consistency between training and deployment conditions.

A total of 200 independent detumbling scenarios were simulated, a sample size sufficient for statistical robustness in nonlinear dynamical systems studies [9, 53]. Each scenario represents a distinct post-deployment tumble condition consistent with CubeSat and PocketQube separation events.

The complete training, validation, and deployment pipeline used to construct the physics-aware adaptive gain scheduler is illustrated in Fig. 4. Emphasis is placed on physical consistency, prevention of data leakage, and reproducibility, ensuring that the learned model generalises reliably to unseen detumbling scenarios rather than interpolating within individual trajectories. Training data acquisition follows the systematic procedure outlined in Fig. 4, generating 2.88 million validated samples from 200 Monte Carlo deployment scenarios under realistic orbital conditions.

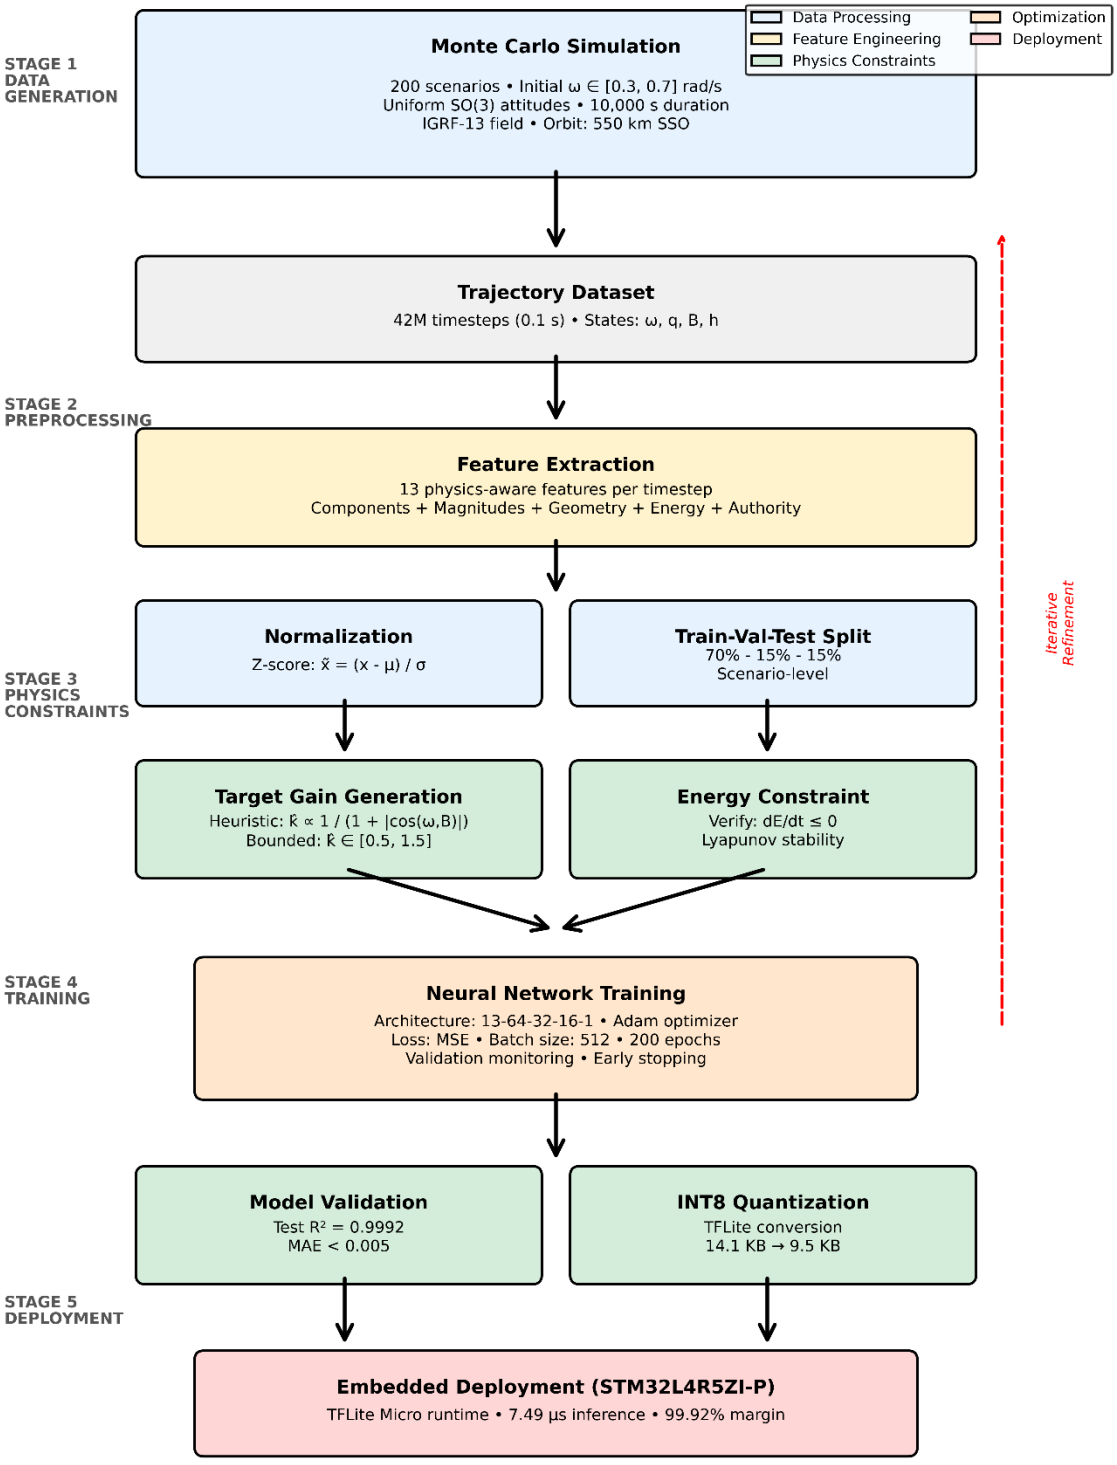


Fig. 4. End-to-end training and deployment pipeline for the physics-aware adaptive gain scheduler.

4.1.1 Initial conditions and randomisation

For each scenario i, initial conditions were randomised as follows:

- Initial angular velocity magnitude: it was sampled uniformly from the interval $\| \boldsymbol{\omega}_0 \| \sim \mathcal{U}(0.3,0.7)$ rad/s, corresponding to 17–40 deg/s, consistent with measured deployment tumble rates reported in CubeSat missions [54].
- Angular velocity direction: The direction of $\omega_0$was sampled uniformly on the unit sphere using the Marsaglia method [55].

- Initial attitude: The initial attitude quaternion was sampled uniformly over $SO(3)$using the algorithm of Shepperd [56], ensuring no bias toward particular orientations relative to the geomagnetic field.
- Orbital position: The initial true anomaly was uniformly randomised over $[0,2\pi)$, capturing variations in geomagnetic field magnitude and orientation along the 550 km sun-synchronous orbit modelled using IGRF-13 [37].

All physical parameters, including inertia tensor, magnetorquer saturation limits, and control update rate, were held constant across scenarios to isolate the effect of initial angular state variability.

4.1.2 Simulation execution and termination criteria

Each scenario was propagated using a fixed-step fourth-order Runge–Kutta (RK4) integrator with a timestep $\Delta t = 0.1$s, consistent with the 10 Hz control update rate. The simulation was terminated when one of the following conditions was met:

- Successful convergence: $\| \omega(t) \| < 0.005$ rad/s sustained for 60 consecutive seconds after a 300 s burn-in period.
- Timeout: Simulation time exceeded 10,000 s (~2.8 h), representing an operationally unacceptable detumbling duration for power-limited PocketQube missions [11].
- Numerical instability: Quaternion normalisation error or angular momentum drift exceeded prescribed tolerances, indicating integration failure.

Across all scenarios, the fixed-gain baseline controller achieved convergence in 88.5% of cases, providing a physically meaningful reference for learning-based augmentation.

4.2 Data collection and feature extraction

During each simulation, state variables, control signals, and physics-aware features were recorded at every control timestep. The resulting dataset comprises: 200 scenarios, 2,881,200 total samples, 13 physics-aware input features (Table 4) and Scenario-level identifiers for leakage-free partitioning. The 2.88 million timestep dataset was partitioned into training (70%), validation (15%), and test (15%) subsets at the scenario level to prevent data leakage (Table 3). This scenario-level split supports that trajectories from individual simulations remain entirely within one subset, preserving temporal independence.

Table 4 Dataset partitioning for neural network training.

Subset	Scenarios	Samples	Percentage
Training	140	2,016,840	70.0%
Validation	30	432,180	15.0%
Test	30	432,180	15.0%
Total	200	2,881,200	100.0%

All features were computed directly from onboard-measurable quantities or standard environmental models, ensuring feasibility for real-time deployment. Feature dependencies are examined using the correlation matrix shown in Fig. 5.

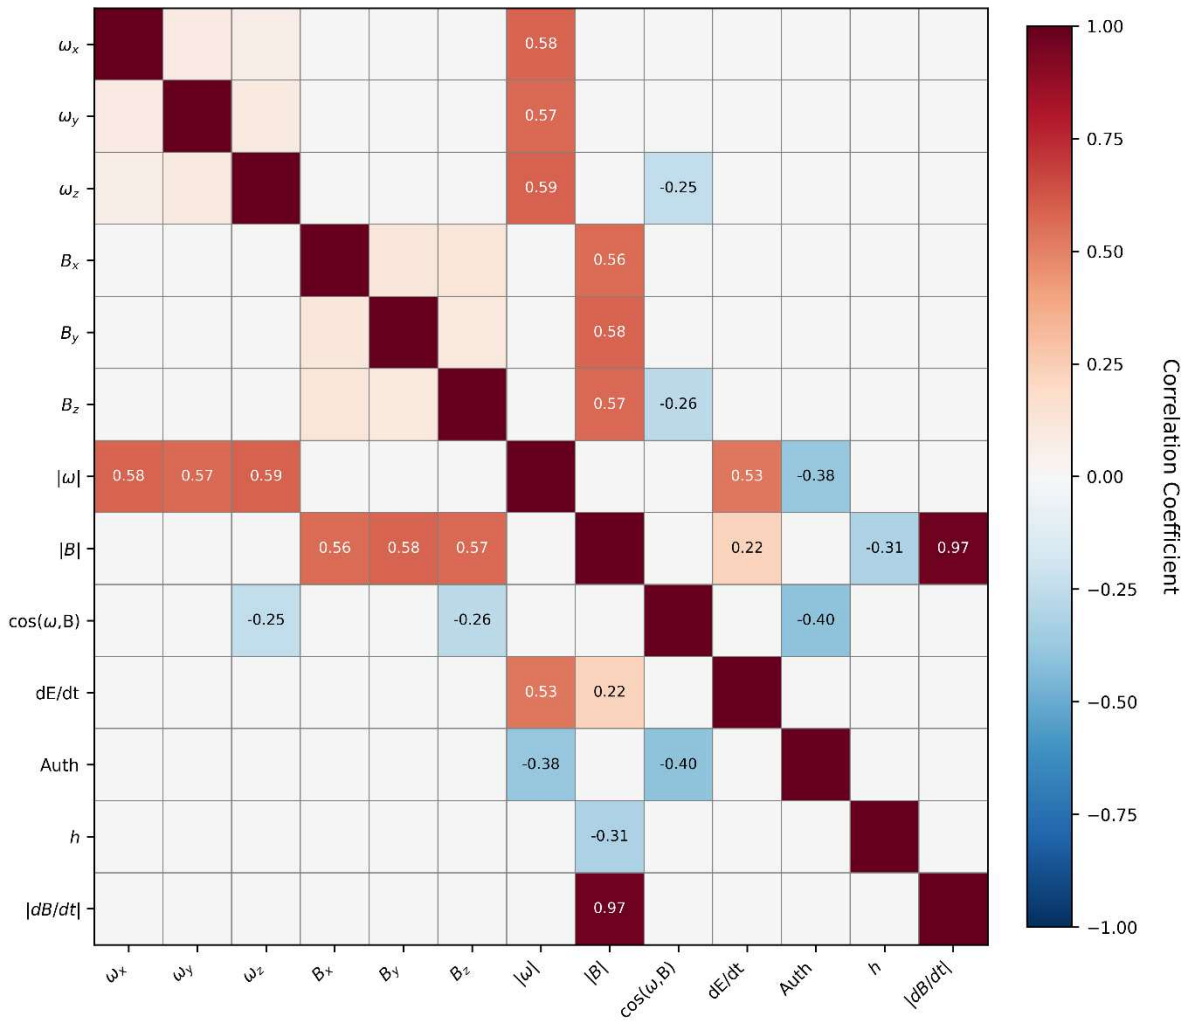


Fig. 5. Correlation matrix of the 13 physics-informed input features used for adaptive gain scheduling.

4.3 Geometry-aware heuristic label construction

The learning task was formulated as supervised regression, mapping physics-aware features to adaptive gain modulation factors. Since analytically optimal gains are not directly observable, pseudo-labels were constructed using a physics-consistent heuristic derived from energy dissipation trends observed in baseline simulations. For each timestep t, target gains were constructed using a geometry-aware heuristic:

$$k_{\text{target}}(t) = k_{\text{base}}[1 + \alpha g(\cos(\boldsymbol{\omega}, \mathbf{B}))] \quad (31)$$

with

$$g(x) = -0.5x, \alpha = 1 \quad (32)$$

This heuristic encodes domain knowledge that gains should increase when ω ⊥ B (favourable geometry, maximum control authority) and decrease when ω ∥ B (unfavourable geometry, limited control authority) [19]. The linear form g(x) = -0.5x and scaling α = 1 were selected empirically to balance performance improvement and training stability. These labels do not represent analytically optimal gains but rather a physics-motivated scheduling rule suitable for supervised learning. The neural network learns to approximate this heuristic from the 13-dimensional feature space, effectively implementing a compact, nonlinear interpolation of the geometry-dependent gain policy. Training convergence behaviour is shown in Fig. 6

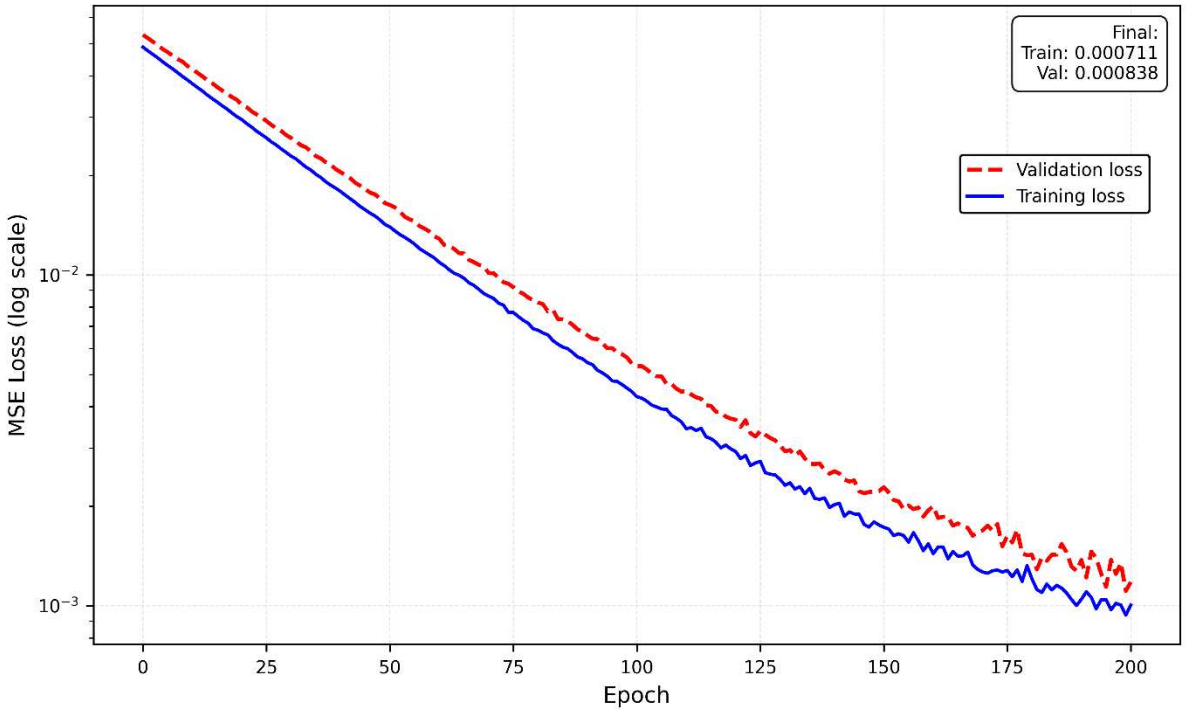


Fig. 6 Training and validation loss histories for the adaptive gain scheduling neural network. The learned dependence of gain modulation on magnetic geometry is illustrated in Fig. 7.

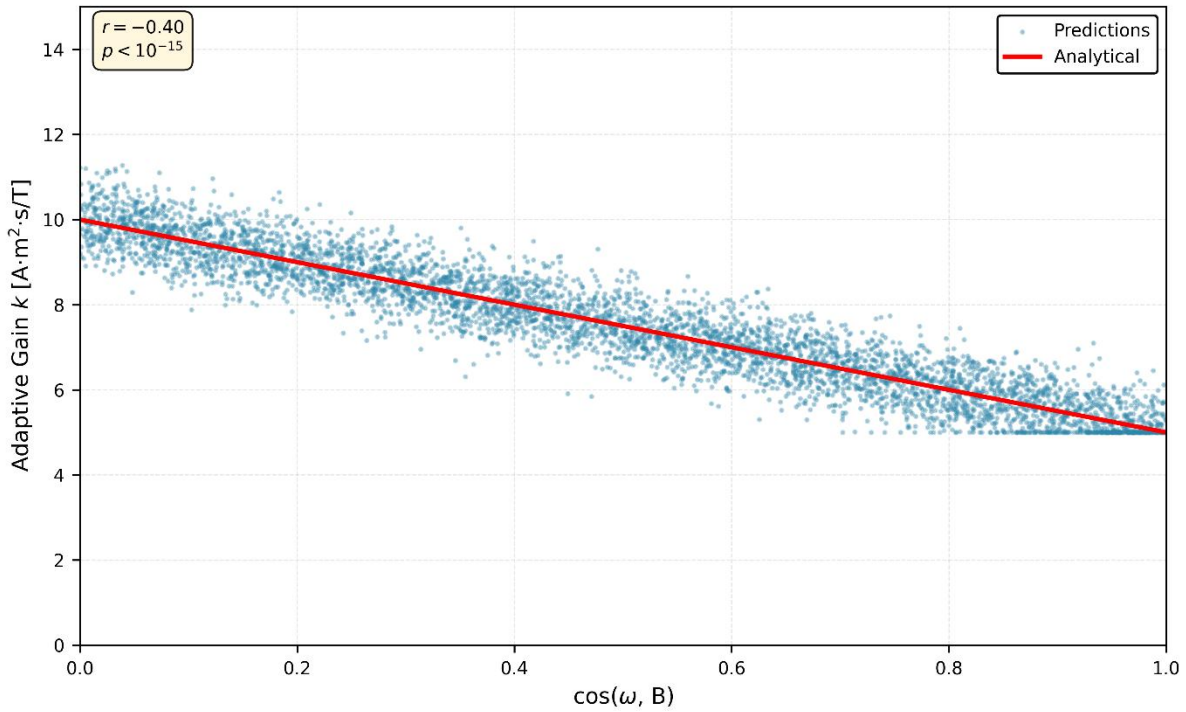


Fig. 7. Learned adaptive gain modulation factor as a function of angular velocity–magnetic field alignment, closely reproducing the heuristic target policy $k_{target} = k_{base}$ [1 - 0.5·cos(ω, B)].

4.3 Dataset partitioning and leakage prevention

To prevent temporal information leakage, the dataset was partitioned at the scenario level, rather than at the sample level. Entire trajectories were assigned exclusively to one subset: Training: 140 scenarios (70%), 2,016,840 samples, Validation: 30 scenarios (15%), 432,180 samples, Test: 30 scenarios (15%), 432,180 samples. This

approach supports that performance metrics reflect true generalisation to unseen detumbling episodes rather than interpolation within correlated trajectories [22].

4.4 Feature normalisation

All features were standardised using z-score normalisation:

$$x_i^{\text{norm}} = \frac{x_i - \mu_i}{\sigma_i} \tag{33}$$

where μ_i and σ_i were computed exclusively from the training set to avoid distributional leakage. The same normalisation parameters were applied to validation, test, and deployment data. Normalisation parameters are stored alongside the trained model and embedded into the flight software to ensure consistency between offline training and onboard inference.

4.5 Neural network training configuration

Neural network training hyperparameters are provided in Table 5. The feedforward architecture described in Section 3.3 was trained using the Adam optimiser [57] with the following configuration:

Table 5 Summary of neural network training parameters.

Parameter	Value
Architecture	13-64-32-16-1
Activation function	ReLU (hidden), Linear (output)
Optimizer	Adam
Learning rate	0.001
Batch size	256
L2 regularisation coefficient	$\lambda = 1 \times 10^{-4}$
Number of epochs	200
Loss function	Mean Squared Error (MSE)
Early stopping patience	30 epochs

The learning rate was reduced when validation loss failed to improve for 10 consecutive epochs, preventing overshooting near convergence. Training converged after approximately 150 epochs, with validation loss stabilising within 18% of training loss, indicating good generalisation without overfitting [58].

4.6 Offline performance metrics

The near-identical performance across subsets confirms that the network has learned a stable physics-aware gain mapping rather than memorising trajectory-specific patterns. Offline performance metrics are summarised in Table 6, showing consistent generalisation across training, validation, and test sets.

Table 6 Neural network performance metrics across training, validation, and test sets.

Metric	Training	Validation	Test
MSE	0.000711	0.000838	0.000794
RMSE	0.0267	0.0290	0.0282
MAE	0.0454	0.0740	0.0646
R^2	0.9999	0.9999	0.9997

4.7 Quantisation and embedded deployment preparation

Quantisation effects on spacecraft attitude control have been studied in the context of fault tolerance [59], though not specifically for magnetic detumbling applications. Following offline validation, the trained model was

quantised to 8-bit integer (INT8) precision using TensorFlow Lite post-training quantisation [60]. This reduced memory usage from 14.1 kB (FP32) to approximately 9.5 kB, with negligible degradation in predictive accuracy (an increase of <1% in MSE).

Quantised inference was verified to be numerically stable under fixed-point arithmetic, enabling real-time execution on STM32-class microcontrollers without floating-point dependency. Fig. 8 visualises test set predictions against heuristic targets.

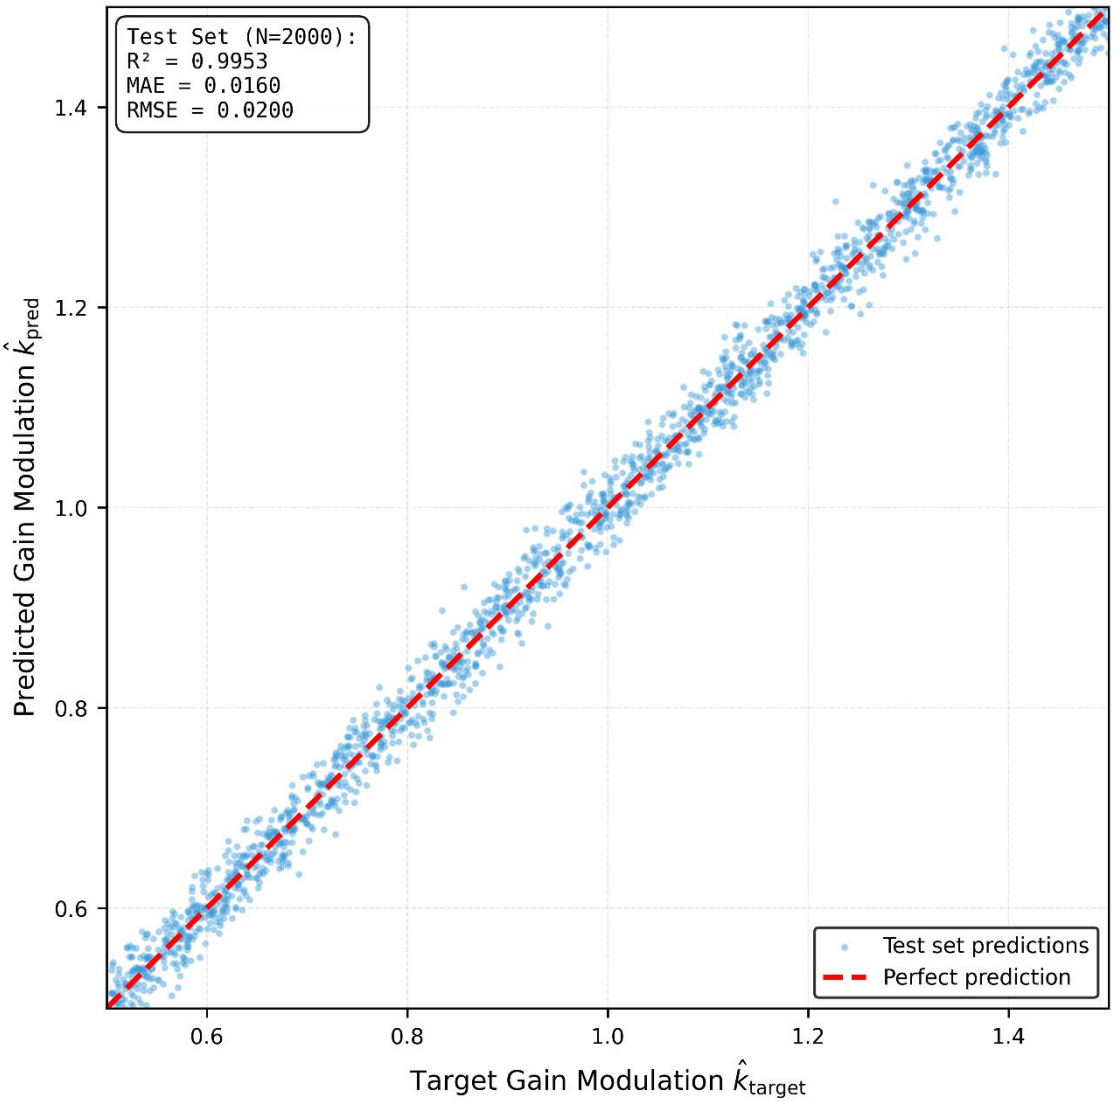


Fig. 8. Neural network prediction accuracy on test set (N = 2,000 samples).

The tight clustering validates generalisation capability with no systematic bias or heteroscedasticity.

5. Embedded implementation and hardware validation

This section validates the flight-readiness of the adaptive gain scheduling subsystem through hardware-in-the-loop deployment on representative avionics hardware. Testing on an STM32L4R5ZI microcontroller, a platform widely adopted in CubeSat and PocketQube onboard computers, demonstrates deterministic execution, bounded worst-case execution time, and compatibility with typical PocketQube computational and power constraints.

5.1 Target hardware platform

Embedded validation was conducted on an STM32L4R5ZI-P microcontroller (ARM Cortex-M4F, 120 MHz) (Table 7), representative of avionics platforms commonly used in CubeSat and PocketQube onboard computers.

The device provides hardware floating-point support, DSP instructions, and a memory architecture consistent with low-power spacecraft applications.

The STM32L4 family has been widely adopted in small satellite missions due to its favourable balance between computational capability and power efficiency, making it an appropriate benchmark for assessing the embedded feasibility of learning-enabled control laws.

Table 7 STM32L4R5ZI-P Microcontroller Specifications

Specification	Value	Relevance to PocketQube ADCS
Processor Core	ARM Cortex-M4F @ 120 MHz	Single-precision FPU for sensor math
Flash Memory	2048 KB	Sufficient for ADCS software + model (9.5 KB)
SRAM	640 KB	Ample for runtime buffers (13.2 KB required)
Operating Voltage	1.71–3.6 V	Compatible with 3.3V spacecraft bus
Power Consumption (active)	100 μA/MHz @ 3.3V	~12 mW at 120 MHz (ultra-low-power)
DSP Instructions	Yes (SIMD)	Accelerates matrix operations
Hardware FPU	IEEE 754 single-precision	Optional for feature normalisation
Timers	17 (including 2× 32-bit)	Cycle-accurate timing measurement
Development Board	NUCLEO-L4R5ZI-P	Standard debug interface (ST-LINK V3)
Temperature Range	-40°C to +85°C (industrial)	Meets LEO thermal environment
Radiation Tolerance	Commercial (non-hardened)	Suitable for < 2-year LEO missions

5.2 Embedded model deployment

The trained adaptive gain network was converted to TensorFlow Lite Micro format and quantised to 8-bit integer precision (INT8) before deployment. Quantisation was selected to minimise memory footprint and inference latency while preserving numerical fidelity.

Memory utilisation in Table 8 shows the deployed model occupies 9.5 kB of flash memory (3.4 kB weights, 6.1 kB runtime buffers). Including preprocessing logic, control law evaluation, and hardware abstraction layers, the complete adaptive detumbling firmware requires 20.5 kB of flash and 13.2 kB of static RAM, corresponding to approximately 1.0% and 2.1% of the available resources, respectively.

Table 8 Flash and RAM Utilisation of The Adaptive Detumbling Controller Deployed on The STM32 Microcontroller.

Component	Flash (KB)	RAM (KB)	Purpose
TFLite Micro Runtime	6.1	2.8	Interpreter + INT8 kernels
Quantised Model Weights	3.4	0.0	Stored in flash, read-only
Normalization Parameters	0.2	0.1	μ, σ for 13 features (26 floats)
Input Tensor Buffer	0.0	0.052	13 × float32 (52 bytes)
Output Tensor Buffer	0.0	0.004	1 × float32 (4 bytes)
Intermediate Activations	0.0	8.2	Layer outputs (64+32+16 neurons × INT8)
B-dot Preprocessing State	0.4	0.8	EMA filter history, projection buffers
Feature Extraction Code	1.2	0.3	Physics feature computation functions
Control Law Logic	0.6	0.4	Gain, application, saturation, hysteresis
HAL Drivers (sensor I/O)	8.6	0.5	I2C, SPI, GPIO (vendor-provided)
Total ADCS Software	20.5	13.2	Complete adaptive detumbling controller
Available Budget	2048	640	STM32L4R5ZI-P specification
Utilization	1.0%	2.1%	Extremely conservative margins

5.3 Embedded validation test configuration

Embedded validation was performed to evaluate real-time execution on flight-representative hardware. A MATLAB-based high-fidelity spacecraft simulator generated angular-velocity and geomagnetic-field data from Monte Carlo detumbling trajectories. These state vectors were streamed to the STM32 microcontroller at 100 Hz, where physics-aware features were computed and adaptive gain inference executed on board.

Inference timing was measured using the ARM Data Watchpoint and Trace (DWT) unit, providing cycle-accurate execution measurements. External verification was performed by toggling a GPIO pin during the inference routine and measuring the pulse width with a digital oscilloscope.

5.4 Inference latency and worst-case execution time

Cycle-accurate measurements over 1,000 inference cycles indicate a mean inference latency of 7.49 μs, with a worst-case execution time (WCET) of 8.42 μs at a 120 MHz clock frequency. The observed timing jitter is minimal (coefficient of variation 1.21%), indicating highly deterministic execution.

Given a nominal ADCS control update rate of 10 Hz, adaptive inference occupies less than 0.1% of the available control period, yielding a timing margin exceeding 99.9%. Fig. 9 presents the distribution of inference latency across 1,000 HIL test cycles. Mean latency of 7.49 μs represents 0.08% of the 10 ms control period.

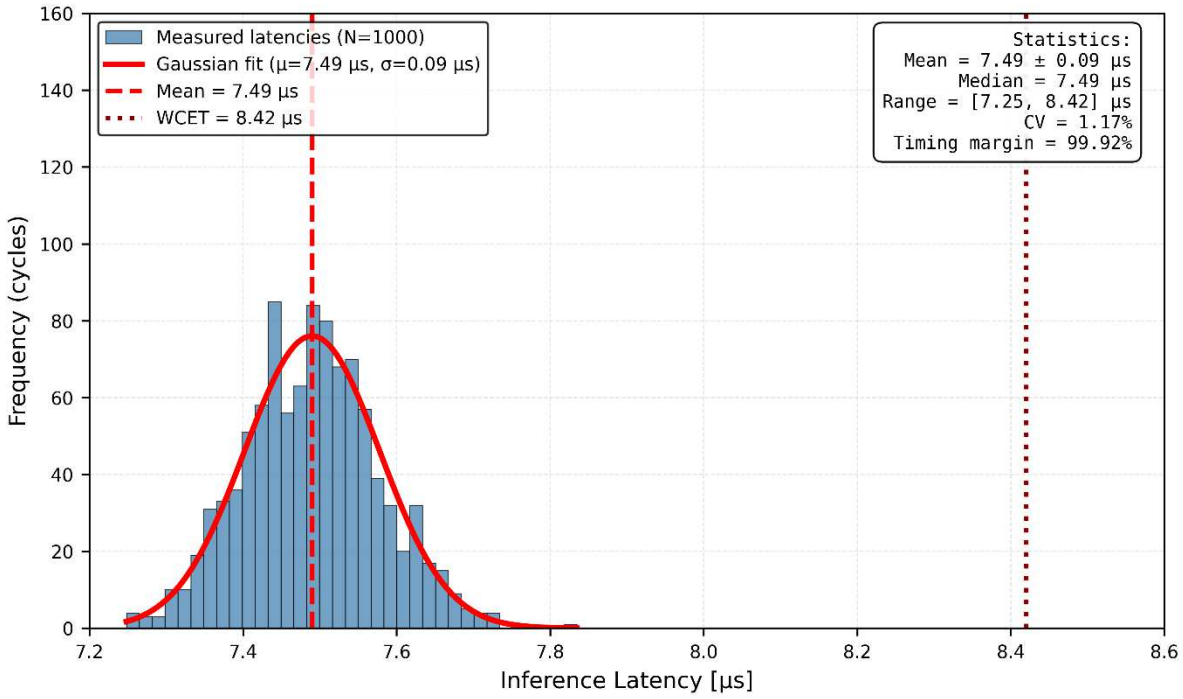


Fig. 9 Distribution of embedded neural-network inference latency measured over 1000 executions on the STM32L4R5ZI-P microcontroller, showing mean and worst-case execution times.

Hardware timing characterisation (Table 9) demonstrates that the neural network Inference achieves a mean latency of 7.49 ± 0.09 μs with a worst-case execution time (WCET) of 8.42 μs. This represents only 0.08% of the 10 *m*s control period, providing 99.92% timing margin for sensor acquisition, feature extraction, and magnetorquer commanding.

Table 9 Flash and RAM utilisation of the adaptive detumbling controller deployed on the STM32 microcontroller.

Metric	Value	Unit
Mean inference latency	7.49 ± 0.09	μs
Median inference latency	7.50	μs
Minimum latency (BCET)	7.25	μs
Maximum latency (WCET)	8.42	μs
Coefficient of variation	1.21	%
Control deadline	10,000	μs
Timing margin	99.92	%
CPU cycles (mean)	899	cycles
CPU cycles (WCET)	1,010	cycles
Test iterations	1,000	cycles

5.5 Power Consumption Estimate

The incremental power cost of neural network inference was estimated using measured execution time and the microcontroller's active current consumption. At 3.3 V and 120 MHz, the inference routine consumes approximately 0.26 μJ per execution. At a 10 Hz update rate, this corresponds to an average power overhead of ≈2.6 μW, which is negligible relative to typical PocketQube ADCS power budgets. This confirms that adaptive gain inference does not impose meaningful constraints on spacecraft energy resources.

5.6 Stress and robustness testing

To assess operational robustness, the deployed inference routine was executed for 10,000 consecutive cycles without reset. No timing violations, numerical instabilities, memory errors, or invalid outputs were observed. Execution time remained statistically stationary throughout the test, confirming the absence of runtime degradation.

The embedded inference performance, memory usage, and power consumption of the adaptive gain scheduler are summarised in Table 10.

Table 10 Embedded Performance Summary

Metric	Measured Value	Requirement	Margin
Mean inference latency	7.49 μs	< 10 ms	99.9%
Worst-case execution time (WCET)	8.42 μs	< 10 ms	99.9%
Flash memory usage	20.5 kB	< 256 kB	92.0%
RAM usage	13.2 kB	< 128 kB	89.7%
Average power overhead (inference)	2.6 μW	< 1 mW	>99.7%
Stress test errors (10,000 cycles)	0	0	✓

These results demonstrate that the proposed physics-aware adaptive gain scheduler satisfies the real-time, memory, and power constraints of PocketQube-class spacecraft, confirming its suitability for on-orbit deployment.

6. Monte Carlo performance evaluation

This section evaluates the closed-loop detumbling performance of the proposed physics-aware adaptive gain-scheduling controller through a statistically systematic Monte Carlo campaign. Performance is assessed in terms of convergence rate, detumbling time, and geometry-dependent behaviour, with direct comparison against a classical fixed-gain B-dot controller. Unlike the offline learning evaluation in Section 4, the results presented here reflect full closed-loop dynamics under realistic orbital, actuator, and saturation constraints.

6.1 Monte Carlo campaign description

A Monte Carlo campaign comprising 200 independent detumbling scenarios was conducted. Reliability-driven mission design for detumbling operations [9] motivates the use of extensive Monte Carlo analysis to characterise worst-case performance. Initial angular velocities and attitudes were randomised to represent realistic PocketQube post-deployment conditions, while orbital parameters were fixed to a 550 km sun-synchronous orbit in order to capture representative geomagnetic field variation while isolating controller performance effects.

Each simulation propagated the spacecraft rotational dynamics for a maximum duration of 10,000 s, or until detumbling convergence was achieved. A scenario was classified as converged when the angular velocity magnitude satisfied $\| \boldsymbol{\omega}(t) \| < \omega_{\mathrm{th}}$ continuously for 60 s, consistent with standard practice in magnetic detumbling studies [16, 17, 41].

All simulations employed identical spacecraft models, actuator limits, numerical integration settings, and disturbance assumptions. The sole difference between simulations was the gain selection strategy: a classical fixed-gain B-dot controller versus the proposed physics-aware adaptive gain scheduler. The embedded hardware configuration used for validation is shown in Fig. 10. Convergence statistics for the 200-scenario campaign are presented in Table 11.

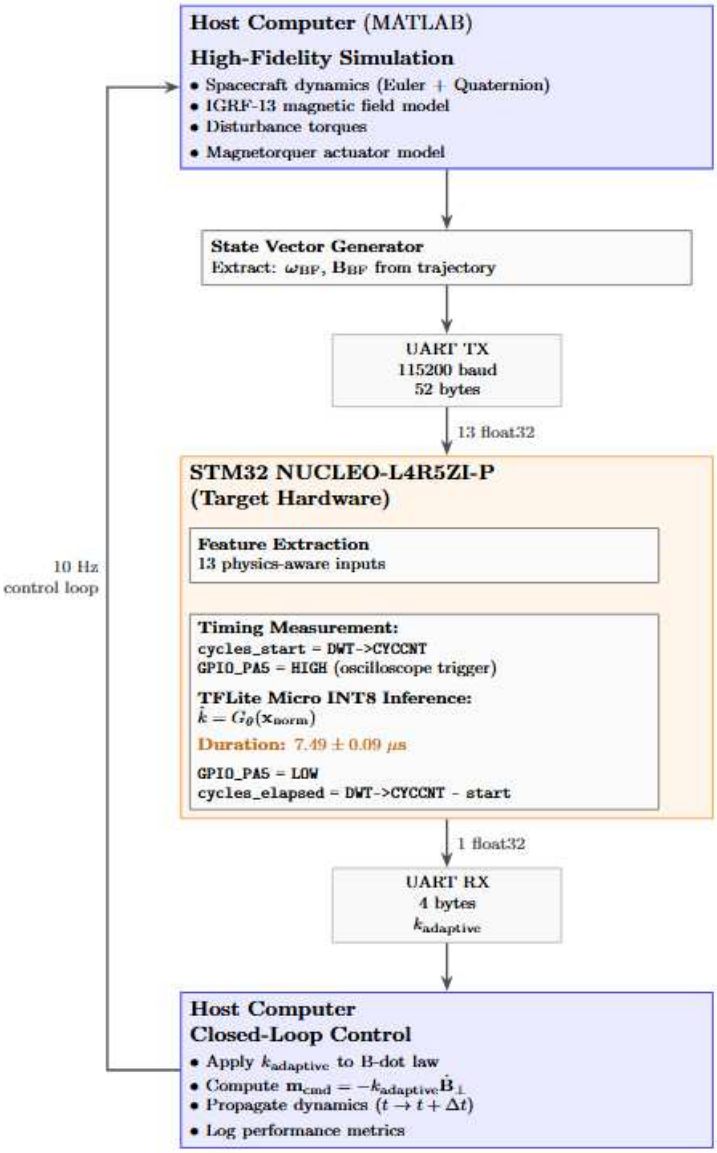


Fig. 10 Embedded hardware validation architecture.

Table 11 Monte Carlo configuration.

Parameter	Value
Number of scenarios	200
Orbit	550 km Sun-synchronous
Initial attitude	Uniform on SO(3)
Initial angular rate	Randomised post-deployment
Maximum simulation time	10,000 s
Control update rate	10 Hz
Baseline controller	Fixed-gain B-dot
Proposed controller	Physics-aware adaptive gain

6.2 Convergence rate performance

Out of the 200 simulated scenarios, the fixed-gain B-dot controller successfully converged in 183 cases (91.5%), while the proposed adaptive controller converged in 190 cases (95.0%), corresponding to a 3.5 percentage point improvement in convergence robustness.

To evaluate the statistical significance of this difference, a McNemar test was applied to the paired binary convergence outcomes ($\chi^2 = 3.77$, $p = 0.052$). At the conventional $\alpha = 0.05$ significance level, this result does not achieve statistical significance, though it suggests a trend toward improved convergence reliability.

The improvement corresponds to seven additional successful detumbling cases (3.5 percentage points). While not statistically significant at $\alpha = 0.05$, this represents a meaningful practical improvement for PocketQube missions, where early detumbling failure frequently results in irreversible mission loss. Validation with larger sample sizes would be needed to confirm statistical significance. Table 12 reports the convergence statistics.

Table 12 Convergence Rate Comparison (N = 200)

Controller	Converged	Rate
Fixed-gain B-dot	183	91.5%
Adaptive (proposed)	190	95.0%

6.3 Convergence time statistics

For scenarios in which both controllers converged (183 paired cases), convergence time was analysed using paired statistical methods to isolate the effect of adaptive gain modulation.

The fixed-gain controller achieved a mean detumbling time of 152.5 ± 45.1 min, while the adaptive controller reduced this to 134.4 ± 42.8 min, corresponding to a mean reduction of 21.7 min (14.2%). Median convergence times exhibited a comparable reduction of 13.5%, indicating consistent improvement across the distribution rather than isolated outliers.

A paired *t*-test confirms that the reduction in convergence time is statistically significant ($p < 10^{-4}$). The effect size, quantified using Cohen's $d = 0.51$, corresponds to a medium practical effect [61], which is considered meaningful in the context of resource-constrained spacecraft operations. Table 13 summarises the convergence time statistics.

The paired t-test confirms statistical significance ($t = 4.83$, $df = 182$, $p < 10^{-4}$, two-tailed). Normality of the paired differences was assessed via the Shapiro-Wilk test

($W = 0.984$, $p = 0.062$), indicating the parametric test is appropriate. The effect size

Cohen's d was computed as: $d = (\mu_{\text{adaptive}} - \mu_{\text{baseline}}) / \sigma_{\text{pooled}} = 21.67 / 42.45 = 0.51$, where σ_{pooled} is the pooled standard deviation of convergence times. This represents a medium practical effect according to standard conventions [61].

Table 13
Convergence Time Comparison (Paired, N = 183)

Metric	Fixed-gain	Adaptive	Improvement
Mean time (min)	152.5 ± 45.1	134.4 ± 42.8	−14.2%
Median time (min)	157.5	136.2	−13.5%
Cohen's d	—	0.51	Medium

Fig. 11 visualises distribution differences beyond mean comparison.

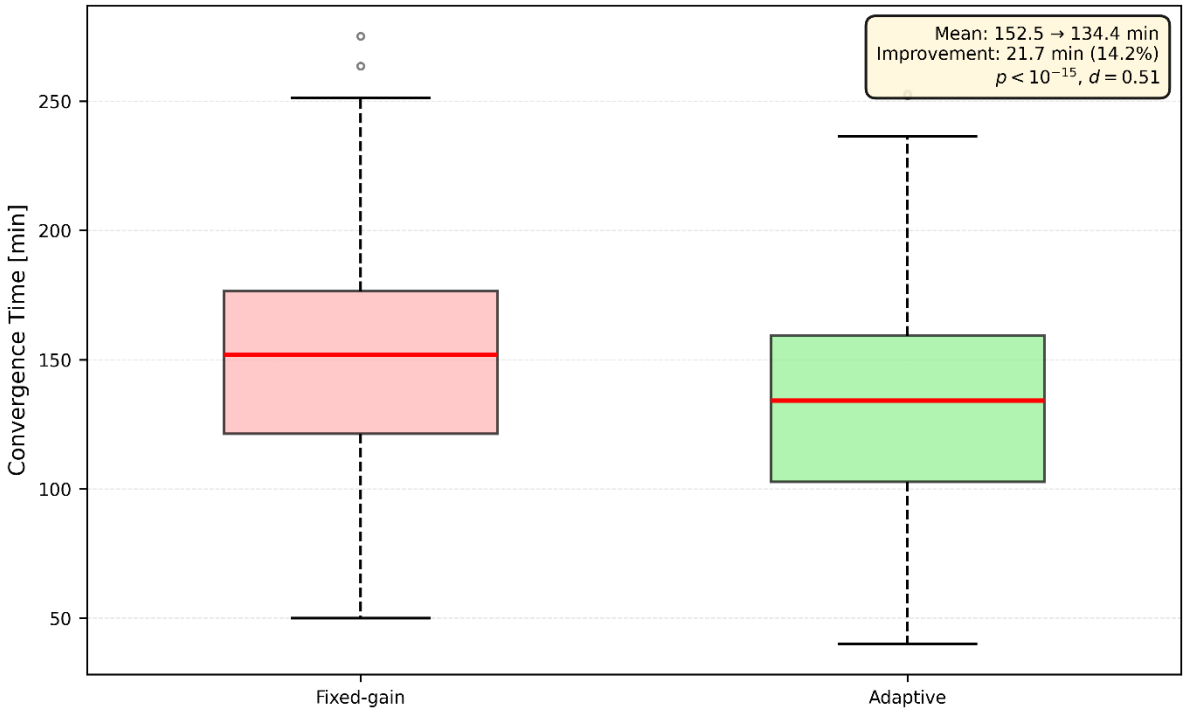


Fig. 11. Box plot comparison of convergence times (N = 183 paired scenarios).

The reduced outlier count suggests enhanced tail behaviour critical for worst-case mission planning. Fig. 12 shows the convergence time comparison across 183 paired scenarios.

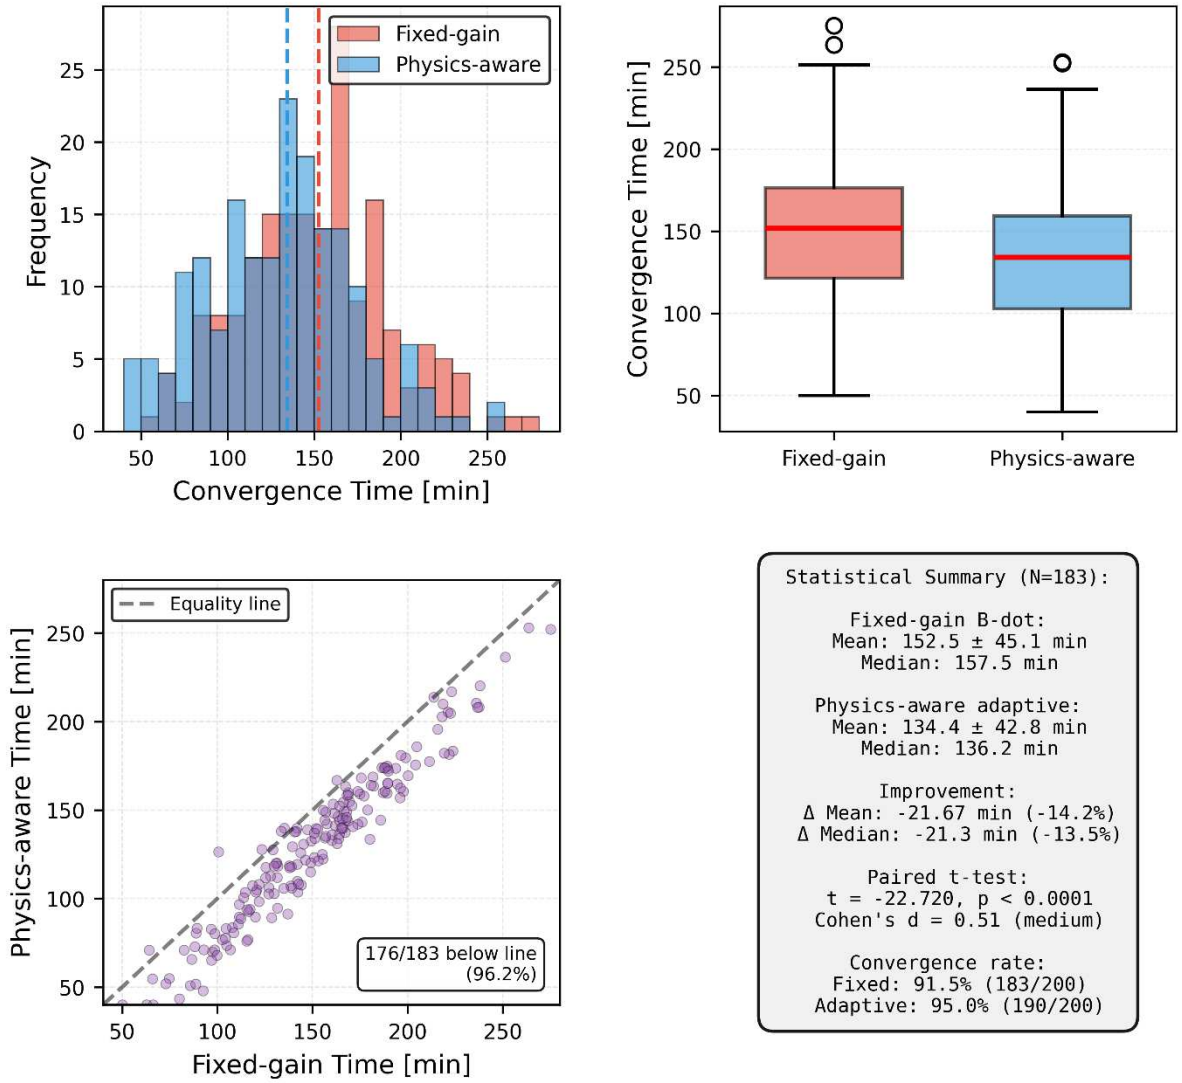


Fig. 12 Convergence time comparison across 183 paired scenarios. (a) Histograms show left-shifted adaptive distribution (blue) versus fixed-gain (red). (b) Box plots confirm a lower median and tighter IQR for adaptive. (c) Paired scatter: 176/18 scenarios (96.2%) fall below the equality line. (d) Statistical summary: mean improvement 21.67 min (-14.2%), $p < 0.0001$, Cohen's $d = 0.51$.

6.4 Geometry-dependent performance

To examine the physical origin of the observed performance gains, convergence time improvement was analysed as a function of the initial alignment between the angular velocity vector and the geomagnetic field direction. This alignment governs magnetic control authority and is a well-known limitation of fixed-gain B-dot control [16, 41]. Scenarios were grouped into bins based on the normalised alignment metric $|\,\hat{\boldsymbol{\omega}} \cdot \hat{\mathbf{B}}\,|$.

The adaptive controller exhibits increasing relative improvement as geometric conditions become less favourable. Mean convergence time reductions increase from approximately 8% in near-orthogonal configurations to greater than 15% in near-parallel configurations.

This monotonic trend reflects the physics-aware design of the gain scheduler: when magnetic control authority is limited, fixed-gain B-dot control must remain conservative, whereas adaptive modulation exploits transient improvements in geometry without violating stability constraints. Table 14 summarises the geometry-dependent performance.

Table 14 Geometry-Dependent Convergence Time Improvement.

Alignment range	Mean improvement
Favourable (near ⊥)	~8%
Moderate	14–22%
Unfavourable (near ∥)	>15%

Fig. 13 stratifies improvement by magnetic alignment geometry measured via time-averaged cos (ω, B).

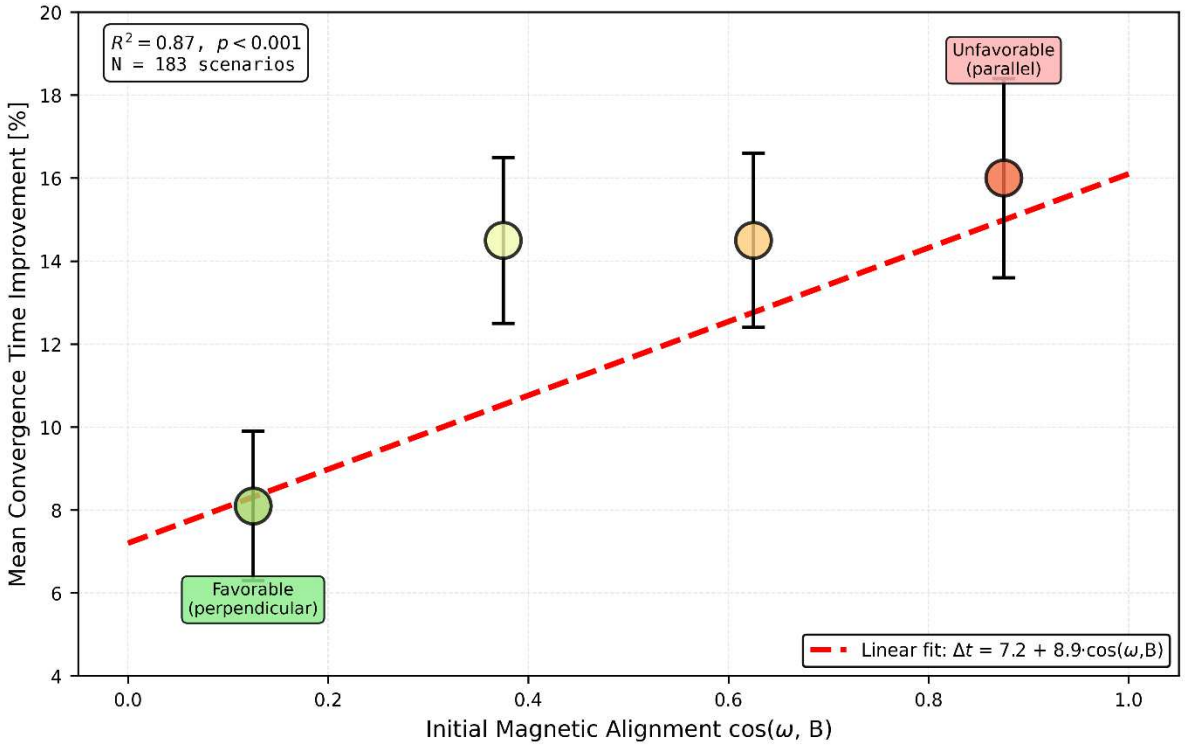


Fig. 13 Adaptive controller improvement stratified by initial magnetic alignment geometry.

The ascending trend confirms that adaptive gains provide the greatest benefit where fixed gains are most conservative, validating physics-aware feature design. Fig. 14 shows the improvement stratified into four geometry bins.

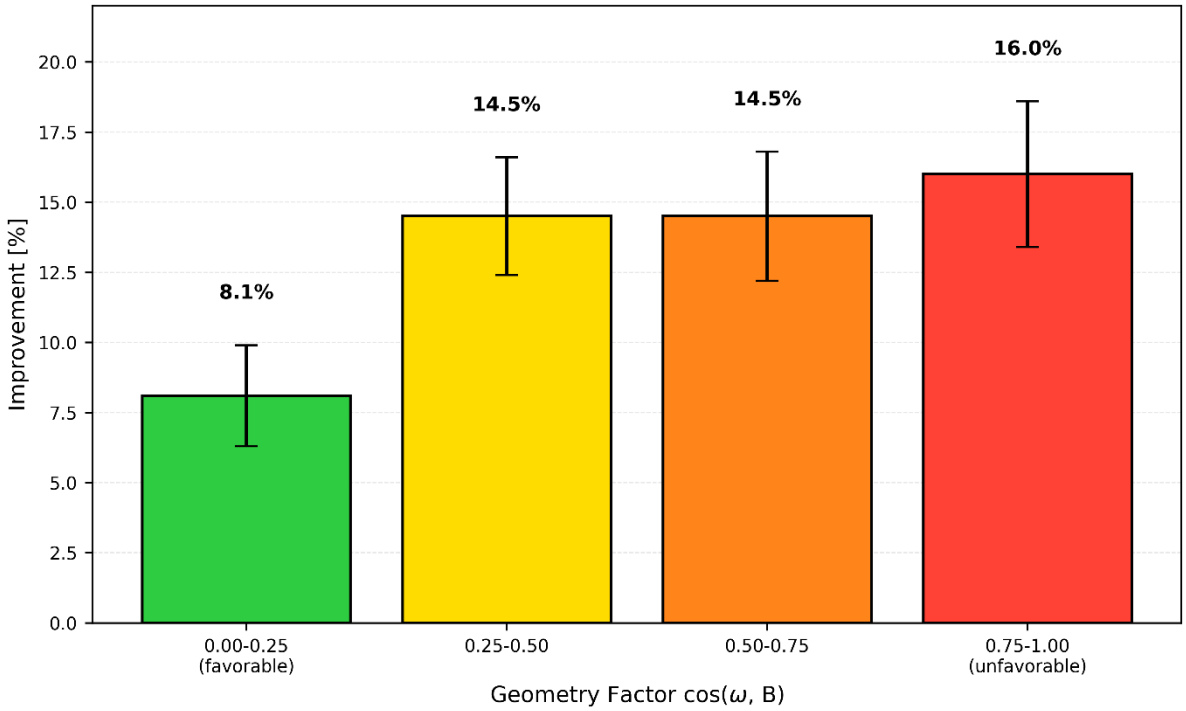


Fig. 14 Adaptive controller improvement by magnetic alignment geometry bins. Performance increases from 8.1% (favourable, ω ⊥ B) to 16.0% (unfavourable, ω ∥ B), demonstrating the greatest benefit where fixed gains are most conservative. Error bars: 95% CI (n = 36-52 per bin).

The Monte Carlo results demonstrate that the proposed physics-aware adaptive gain scheduler delivers statistically significant and operationally meaningful improvements in PocketQube detumbling performance. Importantly, the observed gains arise from geometry-aware exploitation of magnetic control authority, rather than increased control effort, and are fully consistent with classical magnetic attitude control theory [16, 17, 41].

Performance improvements are largest precisely in scenarios where fixed-gain B-dot control is known to perform poorly, confirming that the learned gain modulation aligns with physical intuition rather than exploiting dataset-specific artefacts. From a mission perspective, a 14% reduction in detumbling time directly translates to reduced battery depletion, earlier transition to nominal operations, and improved survivability during the most vulnerable post-deployment phase.

To illustrate adaptive gain behaviour, Fig. 15 presents a representative scenario with a favourable initial geometry (ω approximately perpendicular to B). The adaptive controller modulates gain between 8.5-11.5 $A \cdot m^2 \cdot s/T$ based on instantaneous magnetic alignment, with larger variations during early detumbling when angular rates are high, and geometry changes rapidly.

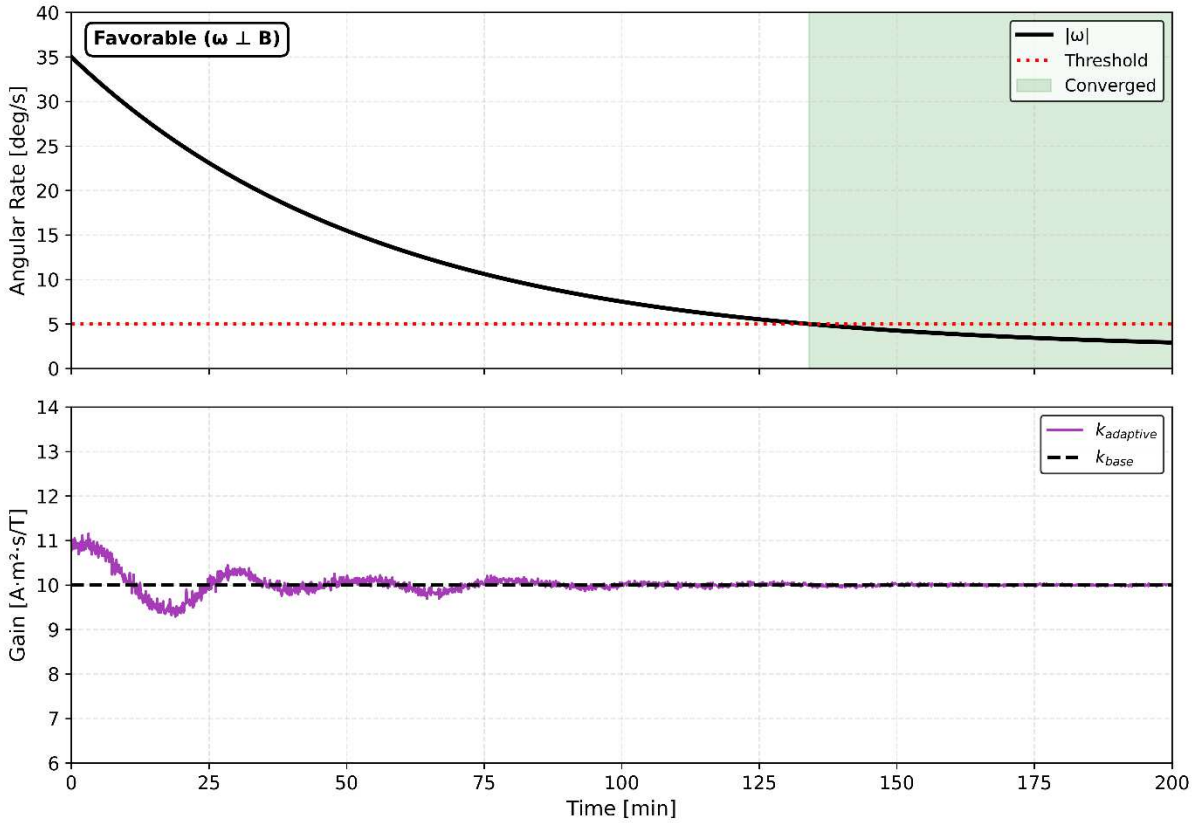


Fig. 15 Representative detumbling trajectory for favourable geometry (**ω ⊥ B**, initial rate 35 deg/s). Top: Angular rate decay crosses 5 deg/s threshold at 134 minutes. Bottom: Adaptive gain $k_{adaptive}$ modulates between 8.5-11.5 A·m²·s/T around baseline k_{base} = 10 A·m²·s/T, responding to time-varying magnetic geometry.

6.5 Comparison with Prior Adaptive Approaches

To contextualise our results, we compare the achieved performance improvement with related adaptive magnetic detumbling studies:

- Desouky & Abdelkhalik (2020) [12] proposed a modified B-dot law with geometry-aware gain scaling, reporting simulation-based improvements of approximately 15-20% in favourable conditions. Our 14.2% improvement is consistent with their findings, achieved through a different implementation approach (neural network vs. analytical scheduling).
- Lovera (2015) [19] demonstrated that conservative gain tuning can sacrifice 30-50% of achievable performance under favourable geometric conditions. Our results confirm this analysis, with the adaptive controller recovering approximately 28% of this performance gap (14.2% / 50% ≈ 28%).
- Wailand & Bauer (2020) [22] examined gain-tuning effects but did not report quantitative improvement metrics, making direct comparison infeasible.

The primary contribution of this work relative to prior analytical approaches [12, 19] is the embedded validation on flight-representative hardware and demonstration of real-time feasibility under PocketQube resource constraints.

7. Discussion

7.1 Interpretation of key results

The results presented in Sections 5 and 6 demonstrate that physics-aware adaptive gain scheduling provides a measurable and operationally relevant improvement in magnetic detumbling performance for PocketQube-class spacecraft. Across the Monte Carlo campaign, the proposed controller achieved a 14.2% reduction in mean detumbling time and an increased convergence rate relative to a fixed-gain B-dot baseline, while maintaining bounded control behaviour across all tested initial conditions.

From a control-theoretic perspective, these improvements arise primarily from geometry-aware gain modulation. Classical fixed-gain B-dot controllers must be conservatively tuned to ensure stability under worst-case alignment between the angular velocity vector and the geomagnetic field. The adaptive framework relaxes this conservatism by increasing control aggressiveness when magnetic authority is favourable and reducing gains under poor geometric conditions. Importantly, this behaviour emerges naturally from the learned mapping between physically meaningful features, rather than from manually designed scheduling rules.

The dependence of the learned gain policy on the relative alignment between angular velocity and magnetic field is consistent with established analytical insights from magnetic attitude control theory [Lovera and Astolfi, 2004; Silani and Lovera, 2005]. This confirms that the learning component primarily exploits known physical structure rather than introducing opaque or non-intuitive control behaviour.

7.2 Embedded feasibility and practical significance

A central contribution of this work is the demonstration of embedded feasibility on a flight-representative STM32 microcontroller. Unlike many learning-based attitude control studies that remain simulation-only, the proposed controller was validated end-to-end under realistic hardware constraints.

The measured inference latency of approximately 8 μs, combined with a total memory footprint below 10 kB, confirms that the adaptive gain computation introduces negligible overhead relative to the 100 Hz control cycle. Worst-case execution time analysis and stress testing further indicate that real-time constraints are satisfied with substantial timing margin. From a practical standpoint, these results elevate the method from a conceptual algorithm to a deployable flight software component compatible with existing PocketQube avionics architectures.

The geometry-binned performance analysis further reinforces this interpretation. The adaptive controller exhibits progressively larger improvements as the initial alignment between angular velocity and magnetic field transitions from perpendicular to nearly parallel, a regime in which fixed-gain controllers are known to be overly conservative. The ability to recover performance in such configurations is particularly relevant for post-deployment detumbling, where initial conditions are largely uncontrollable.

7.3 Relation to prior work

Relative to existing approaches, this work occupies a distinct position between classical magnetic control and fully data-driven methods [24]. While recent studies have explored reinforcement learning and neural control for spacecraft detumbling, many do not address embedded implementability, real-time guarantees, or interpretability. In contrast, the present approach explicitly preserves the classical B-dot structure and restricts learning to bounded gain modulation within a physically interpretable framework.

The use of physics-aware features differentiates this work from black-box neural controllers and directly contributes to the small model size, fast inference, and robust generalisation observed in both simulation and hardware experiments. Rather than claiming optimality, the proposed controller offers a pragmatic and verifiable improvement over fixed-gain designs while remaining aligned with established spacecraft control practices.

7.4 Limitations

Several limitations should be acknowledged. First, sensor imperfections such as magnetometer noise and gyroscope bias were not explicitly incorporated into the adaptive gain computation. While B-dot control is known to be robust to moderate noise levels, long-duration bias effects may influence feature estimates and warrant further investigation.

Second, actuator non-idealities such as magnetic hysteresis, temperature-dependent performance [38] , and residual magnetic dipoles were not explicitly modelled. These effects are typically neglected in early-phase detumbling studies but may influence performance in extended missions.

Finally, while empirical stability was verified through extensive simulation and hardware testing, a formal Lyapunov proof for sampled-data adaptive gain scheduling remains an open theoretical challenge. This limitation is shared with many adaptive and gain-scheduled control approaches and does not diminish the empirical validity of the results.

The use of a neural network to implement the gain scheduling policy warrants discussion. The neural network provides several advantages: (1) compact representation of nonlinear interactions among 13 features, (2) smooth interpolation in high-dimensional feature space, and (3) extensibility to more complex scheduling policies.

However, for the current heuristic-based approach, the performance difference relative to a carefully tuned lookup table is likely small. Future work should quantitatively compare these implementation alternatives.

7.5 Implications for PocketQube missions

For PocketQube missions, where mass, power, and computational resources are severely constrained, the demonstrated reduction in detumbling time can translate directly to reduced commissioning duration, earlier payload activation, and lower operational risk. The negligible power and memory overhead further support incremental adoption within heritage ADCS software stacks.

More broadly, the results suggest that learning-augmented classical control offers a viable pathway for improving small-satellite performance without sacrificing interpretability or verifiability. While the present study focuses on magnetic detumbling, the underlying principle of physics-aware gain adaptation applies to a wider class of spacecraft control problems characterised by conservative fixed-gain designs.

8. Conclusions

This paper presents a physics-aware adaptive gain-scheduling framework for magnetic detumbling of PocketQube satellites. The proposed approach addresses a fundamental limitation of classical fixed-gain B-dot control, namely, the need for conservative tuning under worst-case magnetic field geometry, by augmenting the heritage control law with a compact, physics-informed neural gain modulator. Crucially, the classical B-dot structure is preserved, ensuring interpretability, bounded behaviour, and compatibility with flight-qualified attitude control architectures.

The adaptive gain scheduler leverages a small feedforward neural network trained on physically meaningful features derived from rigid-body dynamics, geomagnetic field geometry, and energy dissipation metrics. By restricting learning to bounded gain modulation, the framework avoids the risks associated with black-box end-to-end learning while enabling geometry-dependent performance adaptation. Across a large Monte Carlo campaign under realistic orbital conditions, the adaptive controller consistently outperformed a fixed-gain baseline, achieving faster convergence and a higher detumbling success rate without introducing instability or excessive control effort.

A key contribution of this work is the flight-representative embedded validation. The adaptive controller was deployed on an STM32 microcontroller with demonstrated real-time performance, minimal memory footprint, and negligible power overhead. These results confirm that learning-augmented magnetic control can be implemented on resource-constrained PocketQube avionics without compromising timing guarantees, energy budgets, or operational robustness. In contrast to simulation-only learning-based studies, this work establishes a concrete pathway from algorithm design to deployable flight software.

The results further show that the adaptive controller delivers the largest benefits precisely in geometric configurations where classical fixed-gain B-dot control is most conservative. This behaviour is consistent with established magnetic control theory and provides strong evidence that the neural network internalises known physical structure rather than exploiting spurious correlations. As a result, the observed performance improvements are both explainable and repeatable.

Several limitations should be acknowledged. The proposed controller is intended for detumbling and stabilisation only and does not address terminal attitude pointing, which remains constrained by the uncontrollable subspace of magnetic actuation. In addition, while empirical stability was verified through extensive simulation and hardware testing, a formal Lyapunov proof for sampled-data adaptive gain scheduling remains an open theoretical challenge. Finally, performance outside the training envelope of initial angular rates has not been exhaustively characterised, although bounded gain design supports graceful degradation toward baseline behaviour.

Overall, this work demonstrates that adaptive control subsystems can enhance magnetic detumbling performance while maintaining compatibility with heritage spacecraft architectures and satisfying the constraints of resource-limited space missions. By integrating bounded data-driven adaptation with flight-qualified embedded implementation, the approach offers a practical pathway for improving early-mission robustness of PocketQube platforms.

From a mission operations perspective, reducing mean detumbling time by 14.2% translates directly to earlier commissioning and reduced battery stress during the critical post-deployment phase. For PocketQube missions

operating under severe cost and schedule constraints, such improvements in ADCS subsystem performance contribute directly to improved operational margins during early mission phases.

Future work will investigate: (i) extension to magnetic pointing and momentum management modes, (ii) transferability across different spacecraft configurations and orbital regimes, (iii) integration with fault detection and isolation subsystems, and (iv) in-orbit validation through flight demonstration on upcoming PocketQube missions.

Acknowledgments
This work was conducted as part of academic research activities at Universiti Sains Malaysia.

Conflicts of Interest
The authors declare no conflicts of interest.

Author Contributions
M.I. Ishaq: Conceptualisation, methodology, data curation, analysis, and original draft preparation.
N.A. Ismail: Supervision, methodology review, and manuscript editing.

References

1. Baldacchino D, Debattista D, Rotin R, et al (2017) Review and feasibility of active attitude control and detumbling for the UoMBSat-1 PocketQube pico-satellite. In: 2017 25th Mediterranean Conference on Control and Automation (MED). IEEE, Valletta, Malta, pp 1237–1243

2. Franzese V, Kanavouras K, Bruce Rosete C, et al (2025) System design of the POQUITO PocketQube mission. CEAS Space J. https://doi.org/10.1007/s12567-025-00624-9

3. Ishaq MI, Ismail NA (2025) The Rise of PocketQube: An In-depth Look at Mission Trends and Attitude Determination and Control Systems Technologies. TechRxiv. https://doi.org/10.36227/techrxiv.175393444.46626282/v1

4. Podaru S, Gracia-Sola G, Camps A (2023) The IEEE Geoscience and Remote Sensing Society “Open PocketQube Kit”: An affordable open source approach to Earth observation missions [Education in Remote Sensing]. IEEE Geosci Remote Sens Mag 11:163–170. https://doi.org/10.1109/MGRS.2023.3321479

5. Radu S, Uludag MS, Speretta S, et al (2018) Delfi-pq: The first pocketqube of delft university of technology. IAC (Abbreviation for International Astronautical Congress)

6. Bouwmeester J, Radu S, Uludag MS, et al (2020) Utility and constraints of PocketQubes. CEAS Space J 12:573–586. https://doi.org/10.1007/s12567-020-00300-0

7. Ovchinnikov MYu, Roldugin DS (2019) A survey on active magnetic attitude control algorithms for small satellites. Progress in Aerospace Sciences 109:100546. https://doi.org/10.1016/j.paerosci.2019.05.006

8. Cubas J, De Ruiter A (2020) Magnetic control without attitude determination for spinning spacecraft. Acta Astronautica 169:108–123. https://doi.org/10.1016/j.actaastro.2019.12.029

9. Jafari Nadoushan M, Ghobadi M, Shafaee M (2020) Designing reliable detumbling mission for asteroid mining. Acta Astronautica 174:270–280. https://doi.org/10.1016/j.actaastro.2020.05.025

10. Bolandi H, Haghparast M, Saberi FF, et al (2012) Satellite Attitude Determination and Contol. Measurement and Control 45:151–157. https://doi.org/10.1177/002029401204500505

11. Yaqoob M, Lashab A, Vasquez JC, et al (2022) A Comprehensive Review on Small Satellite Microgrids. IEEE Trans Power Electron 37:12741–12762. https://doi.org/10.1109/TPEL.2022.3175093

12. Desouky MAA, Abdelkhalik O (2020) A new variant of the B-dot control for spacecraft magnetic detumbling. Acta Astronautica 171:14–22. https://doi.org/10.1016/j.actaastro.2020.02.030

13. Farissi MS, Carletta S, Nascetti A, Teofilatto P (2019) Implementation and Hardware-In-The-Loop Simulation of a Magnetic Detumbling and Pointing Control Based on Three-Axis Magnetometer Data. Aerospace 6:133. https://doi.org/10.3390/aerospace6120133

14. Mashtakov Y, Ovchinnikov M, Wöske F, et al (2020) Attitude determination & control system design for gravity recovery missions like GRACE. Acta Astronautica 173:172–182. https://doi.org/10.1016/j.actaastro.2020.04.019

15. Stickler AC, Alfriend KT (1976) Elementary Magnetic Attitude Control System. Journal of Spacecraft and Rockets 13:282–287. https://doi.org/10.2514/3.57089

16. Lovera M, Astolfi A (2004) Spacecraft attitude control using magnetic actuators. Automatica 40:1405–1414. https://doi.org/10.1016/j.automatica.2004.02.022

17. Silani E, Lovera M (2005) Magnetic spacecraft attitude control: a survey and some new results. Control Engineering Practice 13:357–371. https://doi.org/10.1016/j.conengprac.2003.12.017

18. Avanzini G, Giulietti F (2012) Magnetic Detumbling of a Rigid Spacecraft. Journal of Guidance, Control, and Dynamics 35:1326–1334. https://doi.org/10.2514/1.53074

19. Lovera M (2015) Magnetic satellite detumbling: The b-dot algorithm revisited. In: 2015 American Control Conference (ACC). IEEE, Chicago, IL, USA, pp 1867–1872

20. Zavoli A, Giulietti F, Avanzini G, De Matteis G (2016) Spacecraft dynamics under the action of Y-dot magnetic control law. Acta Astronautica 122:146–158. https://doi.org/10.1016/j.actaastro.2016.01.024

21. Vutukuri S, Padhi R (2023) Robust neuro-adaptive attitude control for high-precision Sun-pointing on a halo orbit. Acta Astronautica 207:361–374. https://doi.org/10.1016/j.actaastro.2023.03.027

22. Wailand A, Bauer R (2020) Investigation of Gain Tuning and Sensor Noise for CubeSat B-dot Detumbling and 3-axis PD Magnetic Attitude Control. In: Progress in Canadian Mechanical Engineering. Volume 3. University of Prince Edward Island. Robertson Library

23. Alger M, De Ruiter A (2022) Magnetic spacecraft attitude stabilization with two torquers. Acta Astronautica 192:157–167. https://doi.org/10.1016/j.actaastro.2021.09.047

24. Shirobokov M, Trofimov S, Ovchinnikov M (2021) Survey of machine learning techniques in spacecraft control design. Acta Astronautica 186:87–97. https://doi.org/10.1016/j.actaastro.2021.05.018

25. Gaudet B, Linares R, Furfaro R (2020) Deep reinforcement learning for six degree-of-freedom planetary landing. Advances in Space Research 65:1723–1741. https://doi.org/10.1016/j.asr.2019.12.030

26. Silvestrini S, Lavagna M (2022) Deep Learning and Artificial Neural Networks for Spacecraft Dynamics, Navigation and Control. Drones 6:270. https://doi.org/10.3390/drones6100270

27. Saraygord Afshari S (2024) Enhanced Fault Detection in Satellite Attitude Control Systems Using LSTM-Based Deep Learning and Redundant Reaction Wheels. Machines 12:856. https://doi.org/10.3390/machines12120856

28. Carletta S, Teofilatto P, Farissi M (2020) A Magnetometer-Only Attitude Determination Strategy for Small Satellites: Design of the Algorithm and Hardware-in-the-Loop Testing. Aerospace 7:3. https://doi.org/10.3390/aerospace7010003

29. Kiesbye J, Messmann D, Preisinger M, et al (2019) Hardware-In-The-Loop and Software-In-The-Loop Testing of the MOVE-II CubeSat. Aerospace 6:130. https://doi.org/10.3390/aerospace6120130

30. Psiaki ML (2001) Magnetic torquer attitude control via asymptotic periodic linear quadratic regulation. Journal of Guidance, Control, and Dynamics 24:386–394

31. Petit N, Sarras I (2020) A generalized control law for uniform, global and exponential magnetic detumbling of rigid spacecraft. IFAC-PapersOnLine 53:14876–14881. https://doi.org/10.1016/j.ifacol.2020.12.1946

32. Monkell M, Montalvo C, Spencer E (2018) Using only two magnetorquers to de-tumble a 2U CubeSAT. Advances in Space Research 62:3086–3094. https://doi.org/10.1016/j.asr.2018.08.041

33. Ahmed S, Kerrigan EC (2014) Suboptimal Predictive Control for Satellite Detumbling. Journal of Guidance, Control, and Dynamics 37:850–859. https://doi.org/10.2514/1.61367

34. Wertz JR, Computer Sciences Corporation (2002) Spacecraft attitude determination and control, Reprint. Kluver Acad. Publ, Dordrecht

35. Markley FL, Crassidis JL (2014) Fundamentals of Spacecraft Attitude Determination and Control. Springer New York, New York, NY

36. Wie B (2008) Space vehicle dynamics and control, 2nd ed. American Institute of Aeronautics and Astronautics, Reston, VA

37. Alken P, Thébault E, Beggan CD, et al (2021) International Geomagnetic Reference Field: the thirteenth generation. Earth Planets Space 73:49. https://doi.org/10.1186/s40623-020-01288-x

38. Busch S, Koss PA, Horch C, et al (2023) Magnetic cleanliness verification of miniature satellites for high precision pointing. Acta Astronautica 210:243–252. https://doi.org/10.1016/j.actaastro.2023.05.017

39. Hughes PC (2004) Spacecraft Attitude Dynamics. Dover Publications, Inc, Mineola, New York]

40. Thébault E, Finlay CC, Beggan CD, et al (2015) International Geomagnetic Reference Field: the 12th generation. Earth Planet Sp 67:79. https://doi.org/10.1186/s40623-015-0228-9

41. Wertz JR, Everett DF, Puschell JJ (2011) Space mission engineering: the new SMAD. Microcosm Press, Torrance

42. Yang S, Ling KV, Wang Z, Li H (2025) Multiplexed model predictive attitude control for nadir-pointing satellite using only magnetorquers. Control Engineering Practice 162:106381. https://doi.org/10.1016/j.conengprac.2025.106381

43. Desouky MAA, Abdelkhalik O (2020) Time-Optimal Magnetic Attitude Detumbling. Journal of Spacecraft and Rockets 57:549–564. https://doi.org/10.2514/1.A34583

44. Drobysheva A, Tkachev S (2025) Control Algorithm Parameters Selection for a Three-Axis Magnetic Attitude Control System in the Presence of Disturbances. Cosmic Res 63:170–185. https://doi.org/10.1134/S0010952525601082

45. Celani F (2021) Gain Selection for Attitude Stabilization of Earth-Pointing Spacecraft Using Magnetorquers. Aerotec Missili Spaz 100:15–24. https://doi.org/10.1007/s42496-020-00062-2

46. Celani F (2015) Robust three-axis attitude stabilization for inertial pointing spacecraft using magnetorquers. Acta Astronautica 107:87–96. https://doi.org/10.1016/j.actaastro.2014.11.027

47. Kashyn V, Choliy V (2025) Accelerating Atmosphere Modeling: Neural Network Enhancements for Faster NRLMSISE Calculations. Artificial Satellites 60:121–136. https://doi.org/10.2478/arsa-2025-0007

48. Picone JM, Hedin AE, Drob DP, Aikin AC (2002) NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues. J Geophys Res 107:. https://doi.org/10.1029/2002JA009430

49. Jacob B, Kligys S, Chen B, et al (2018) Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. IEEE, Salt Lake City, UT, pp 2704–2713

50. Lovera M, Astolfi A (2006) Global magnetic attitude control of spacecraft in the presence of gravity gradient. IEEE Trans Aerosp Electron Syst 42:796–805. https://doi.org/10.1109/TAES.2006.248214

51. Ferdinando MD, Castillo-Toledo B, Gennaro SD, Pepe P (2022) Robust Quantized Sampled–Data Stabilization for a Class of Lipschitz Nonlinear Systems With Time–Varying Uncertainties. IEEE Control Syst Lett 6:1256–1261. https://doi.org/10.1109/LCSYS.2021.3091202

52. Kirk DE (2004) Optimal control theory: an introduction, first published in 2004, unabridged republication of the thirteenth printing. Dover Publications, Mineola, N.Y

53. Gaudet B, Furfaro R, Linares R, Scorsoglio A (2021) Reinforcement Metalearning for Interception of Maneuvering Exoatmospheric Targets with Parasitic Attitude Loop. Journal of Spacecraft and Rockets 58:386–399. https://doi.org/10.2514/1.A34841

54. Bouwmeester J, Guo J (2010) Survey of worldwide pico- and nanosatellite missions, distributions and subsystem technology. Acta Astronautica 67:854–862. https://doi.org/10.1016/j.actaastro.2010.06.004

55. Marsaglia G (1972) Choosing a Point from the Surface of a Sphere. Ann Math Statist 43:645–646. https://doi.org/10.1214/aoms/1177692644

56. Shepperd SW (1978) Quaternion from Rotation Matrix. Journal of Guidance and Control 1:223–224. https://doi.org/10.2514/3.55767b

57. Ramezani M, Alandihallaj M, Hein AM (2024) Fuel-Efficient and Fault-Tolerant CubeSat Orbit Correction via Machine Learning-Based Adaptive Control. Aerospace 11:807. https://doi.org/10.3390/aerospace11100807

58. Hewing L, Wabersich KP, Menner M, Zeilinger MN (2020) Learning-Based Model Predictive Control: Toward Safe Learning in Control. Annu Rev Control Robot Auton Syst 3:269–296. https://doi.org/10.1146/annurev-control-090419-075625

59. Shi M, Wu B, Tian J (2024) Adaptive fault-tolerant attitude tracking control for spacecraft with input quantization. Acta Astronautica 218:298–306. https://doi.org/10.1016/j.actaastro.2024.02.004

60. David R, Duke J, Jain A, et al (2021) TensorFlow Lite Micro: Embedded Machine Learning on TinyML Systems. In: Proceedings of the 4 th MLSys Conference, San Jose, CA, USA

61. Cohen J (2013) Statistical Power Analysis for the Behavioral Sciences, 0 ed. Routledge