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.

Air-Bearing Testbeds for CubeSat and Nanosatellite Attitude Determination and Control Systems: A Systematic Review of Hardware, Control, and Validation Methodologies

Islam Ibrahim
Zagazig University
Mohamed Riad Ghazy
Zagazig University
Aboubakr M. Elhady
abobakr.elhady@ecu.edu.eg

Egyptian Chinese University

Research Article

Keywords: CubeSat, nanosatellite, attitude determination and control system, air-bearing testbed, hardware-in-the-loop testing, sensor fusion, adaptive control, systematic review

Posted Date: September 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10893767/v1

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

Additional Declarations: No competing interests reported.

Air-Bearing Testbeds for CubeSat and Nanosatellite Attitude Determination and Control Systems: A Systematic Review of Hardware, Control, and Validation Methodologies

Islam Ibrahim[1], Mohamed Riad Ghazy[1] , Aboubakr M. Elhady[2,*]
[1]Faculty of Engineering, Zagazig University, Zagazig, Egypt
[2]Faculty of Engineering and Technology, Egyptian Chinese University, Cairo, Egypt
*Corresponding author: abobakr.elhady@ecu.edu.eg

Abstract

Attitude Determination and Control Systems (ADCS) govern spacecraft pointing, stability, and mission-critical orientation, while attitude-control anomalies are among the frequently reported contributors to on-orbit spacecraft failure. Air-bearing testbeds provide a near-frictionless, low-torque rotational environment that approximates on-orbit dynamics, enabling hardware-in-the-loop (HIL) verification of sensors, actuators, and control algorithms before launch. However, published surveys remain fragmented across hardware, actuator, and control domains. This review addresses this gap through a systematic synthesis of 40 peer-reviewed and technical studies published between 2003 and 2025, identified from an initial 131 records across Scopus, Web of Science, IEEE Xplore, the AIAA Aerospace Research Center, and Google Scholar. Studies were classified across ten analytical dimensions covering hardware, structure, balancing, actuation, sensing and fusion, control, disturbance modelling, validation, AI/ML, and historical evolution. The synthesis reveals a transition from costly, institution-specific systems toward modular, open-source, and additively manufactured testbeds, alongside increasing adoption of adaptive and robust control. Key gaps remain in sensor-fusion validation, thruster-based actuation, and learning-based control. Six comparative tables benchmark 17 representative testbeds and compare major technologies. An order-of-magnitude analysis shows that a 10 μm centre-of-mass/centre-of-rotation offset on a 5 kg platform produces approximately 0.5 mN·m gravitational disturbance torque, highlighting the ground-to-orbit reality gap. The review identifies higher-fidelity disturbance compensation, standardized metrics, digital twins, and AI-assisted diagnostics as key priorities.

Keywords: CubeSat; nanosatellite; attitude determination and control system; air-bearing testbed; hardware-in-the-loop testing; sensor fusion; adaptive control; systematic review

1. Introduction

1.1 Background and motivation

The past two decades have transformed access to space. Universities, research institutions, and commercial developers now routinely design, build, and operate CubeSats and nanosatellites at a fraction of the cost and development time of conventional missions. This democratisation has been accompanied by rising mission ambition from simple technology demonstrators to Earth-observation, communication, and formation-flying platforms placing increasing demands on the subsystems responsible for spacecraft pointing and stability.

The Attitude Determination and Control System (ADCS) establishes and maintains spacecraft orientation throughout a mission. Accurate attitude control underpins communication link budgets, payload pointing, solar-power generation, and thermal management. Analyses of on-orbit anomalies consistently rank attitude-control and related avionics faults among the leading contributors to degraded or lost missions [1], and small-satellite programmes are particularly exposed because of their compressed qualification schedules. Yet ADCS design and validation remain difficult, constrained simultaneously by strict mass

properties, volume, power, and cost budgets and by the largely unobservable conditions of the space environment.

Air-bearing testbeds have consequently become the principal ground-based platform for ADCS validation, because they reproduce a near-frictionless, low-torque rotational environment that approximates the rotational dynamics encountered in orbit [2], [3], [4]. By supporting hardware-in-the-loop (HIL) testing, these facilities allow engineers to exercise flight sensors, actuators, and control software under dynamically representative conditions prior to launch, reducing development risk and cost. It should be stated explicitly at the outset that an air bearing emulates only the *rotational* dynamics of a torque-free body: translational microgravity, orbital environment loading, vacuum, and thermal cycling are not reproduced, and residual gravitational, frictional, aerodynamic, and cable torques persist. The fidelity limits that follow from this are quantified in Section 3.6.

1.2 Related reviews and research gap

Air-bearing spacecraft simulators were first surveyed comprehensively by Schwartz et al. [2], whose historical review of planar, spherical, and tabletop simulators remains the standard reference for the pre-CubeSat era. Subsequent work has documented individual facilities or specific technical aspects of CubeSat ADCS testing — magnetic-field simulation using Helmholtz and Merritt coil cages [5], [6], automatic mass-balancing design and control [7], [8], [9], and low-cost or open-source construction methods [10], [11], [12], [13]. A recent review has surveyed CubeSat ADCS testbeds with an emphasis on magnetic, optical, and simulation-in-the-loop/hardware-in-the-loop (SIL/HIL) approaches [14], and dedicated verification-and-validation facilities continue to be reported at the institutional level [15], [16].

Three deficits remain. First, no existing survey traces the two-decade transition of air-bearing hardware from early government-scale systems to democratised, additively manufactured platforms while retaining a consistent basis for technical comparison. Second, no survey applies a single classification scheme spanning hardware, actuation, sensing, control, disturbance modelling, validation methodology, and AI/ML integration; the recent review of [14] is organised primarily by simulation environment rather than by testbed capability, and does not treat learning-based control or quantitative disturbance budgeting. Third, cross-study comparison is impeded by the absence of a structured benchmark relating reported testbed capability to explicit, stated limitations. This fragmentation makes it difficult for a new research group to determine, from the literature alone, which design choices are mature, which remain open, and where AI/ML and digital-twin methods might most usefully be introduced. The present review is intended to close these three deficits.

1.3 Contributions

This paper makes the following contributions:

- A systematic, PRISMA-informed search and screening process reducing an initial pool of 131 records to 40 studies meeting explicit inclusion criteria for air-bearing ADCS testbed research (Section 2).
- A ten-dimension classification framework — hardware configuration, structural design, balancing, actuation, sensing and fusion, control algorithms, disturbance modelling, validation methodology, AI/ML integration, and historical evolution — applied consistently across the corpus (Section 3).
- A comparative benchmarking table of seventeen representative testbeds reporting method, key features, limitations, and relation to the present synthesis (Table 1).
- Five thematic comparison tables covering actuator technologies (Table 2), sensor-fusion methods (Table 3), control algorithms (Table 4), disturbance sources and compensation (Table 5), and AI/ML applications (Table 6), each reporting validation level alongside capability so that simulated and experimentally demonstrated results are not conflated.

- A quantitative disturbance-budget analysis (Section 3.6, Eqs. (4)–(6)) relating achievable balancing precision to residual torque and to available actuator authority, providing a physical basis for the reality-gap discussion.
- Quantitative synthesis of configuration and actuator-technology trends across the 40 reviewed studies (Fig. 5).
- Explicit identification of underexplored areas — sensor-fusion validation, thruster testing, and learning-based control — with a structured research agenda (Section 4).

1.4 Paper organisation

Section 2 describes the review methodology, including search strategy, inclusion criteria, classification framework, and the analytical model used for disturbance budgeting. Section 3 presents results and discussion, organised by historical evolution, hardware design, actuator technologies, sensor integration, control algorithms, disturbance modelling, validation methodology, AI/ML applications, and comparative synthesis. Section 4 discusses open challenges and research directions. Section 5 concludes.

2. Methodology

This review follows a systematic literature review protocol adapted from the PRISMA 2020 statement for identification, screening, eligibility assessment, and inclusion (Fig. 1).

2.1 Search strategy and databases

Records were retrieved from five sources — Scopus, Web of Science, IEEE Xplore, the AIAA Aerospace Research Center, and Google Scholar. The search was executed on Jan, 2026 and covered publications from 2003 to mid-2025. The Boolean search string applied, with field limits and adaptations per database recorded, was:

> (“air bearing” OR “air-bearing” OR “spherical air bearing” OR “attitude simulator”) AND (CubeSat OR nanosatellite OR “small satellite”) AND (“attitude determination” OR “attitude control” OR ADCS) AND (“hardware-in-the-loop” OR HIL OR testbed OR “test bed”)

The initial search returned 131 unique records after 60 duplicates were removed.

2.2 Inclusion and exclusion criteria

Records were screened at title/abstract level to remove studies unrelated to air-bearing hardware, for example purely numerical ADCS studies with no physical testbed. Remaining records underwent full-text eligibility assessment against three criteria: (i) the study describes a physical air-bearing testbed (spherical, hemispherical, or planar) for CubeSat- or nanosatellite-class ADCS; (ii) the study reports experimental validation rather than proposal-only or purely conceptual design; and (iii) the study is available as a peer-reviewed article, conference paper, or institutionally archived thesis or technical report containing sufficient technical detail for classification across the ten dimensions of Section 2.3. Records were excluded where the platform was a translational or free-floating planar robot without rotational attitude emulation, where the satellite class exceeded 50 kg, or where the full text was unavailable. In total, across deduplication, title/abstract screening, and full-text eligibility assessment, 91 of the 131 originally identified records were excluded (Fig. 1), yielding the final corpus of 40 studies.

Screening was performed independently by 3 reviewers; disagreements were resolved by discussion, and inter-rater agreement at the full-text stage was Cohen’s $\kappa = 0.75$. Applying these criteria yielded 40 studies for detailed synthesis, spanning 2003–2025 (Fig. 1). The reference list of this review is larger than the corpus, because contextual sources — competing reviews, spacecraft-failure statistics, and in-orbit demonstrations are cited but are not members of the reviewed corpus.

Fig. 1. PRISMA 2020 flow diagram for identification, screening, eligibility, and inclusion of the 40 studies synthesized in this review. Records identified, n = 131; duplicates removed before screening, n = 60; records screened at title/abstract, n = 71; records excluded at title/abstract, n = 30; reports assessed for eligibility at full text, n = 41; reports excluded with reasons, n = 1; studies included in structured synthesis, n = 40.

2.3 Classification framework

Each included study (included n = 40) was coded across ten records included analytical dimensions selected to span the testbed development lifecycle: (1) air-bearing hardware configuration and structural design; (2) balancing mechanism; (3) actuator technology — reaction wheels, magnetorquers, control moment gyroscopes (CMGs), thrusters; (4) sensor integration and fusion; (5) control algorithm class — classical, nonlinear, adaptive/robust, learning-based; (6) disturbance modelling and compensation; (7) validation methodology, including HIL, model-in-the-loop (MIL), maneuver types, performance metrics, and ground-truth measurement; (8) AI/ML application; (9) chronological positioning; and (10) reported cost and accessibility.

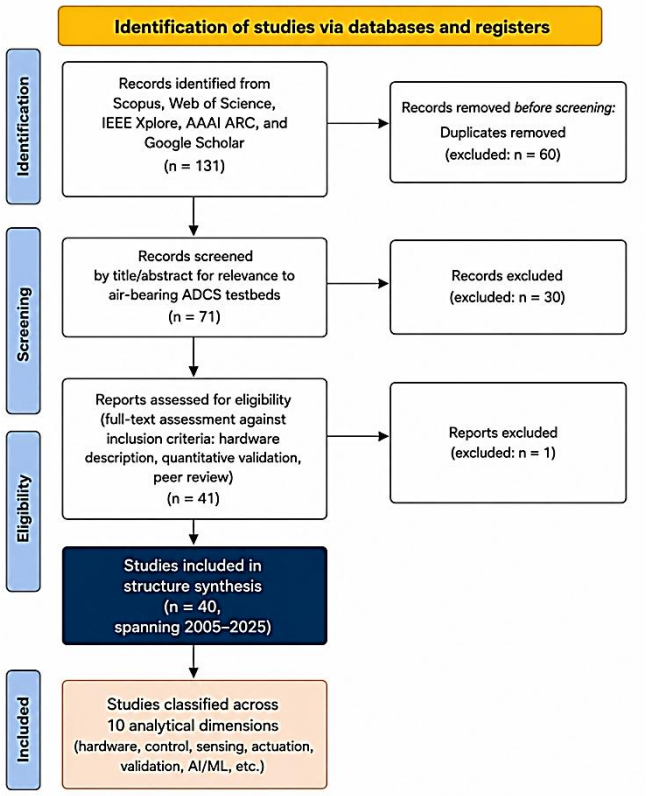


Fig. 1:PRISMA Diagram

2.4 Analytical framework for disturbance budgeting

To place the reviewed hardware on a common quantitative footing, the following standard rigid-body relations are used throughout Section 3. The rotational dynamics of a platform carrying momentum-exchange actuators are

$$\mathbf{J}\dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (\mathbf{J}\boldsymbol{\omega} + \mathbf{h}_w) = \boldsymbol{\tau}_c + \boldsymbol{\tau}_d \tag{1}$$

where $\mathbf{J}$ (kg·m²) is the platform mass inertia tensor, $\boldsymbol{\omega}$ (rad·s⁻¹) the body angular-velocity vector, $\mathbf{h}_w$ (N·m·s) the total wheel angular momentum, $\boldsymbol{\tau}_c$ (N·m) the commanded control torque, and $\boldsymbol{\tau}_d$ (N·m) the net disturbance torque. Attitude is propagated using the quaternion kinematic relation

$$\dot{\mathbf{q}} = 1/2\ \boldsymbol{\Omega}(\boldsymbol{\omega})\,\mathbf{q} \tag{2}$$

with $\mathbf{q} = [q_0,\ \mathbf{q}_v^{\mathrm{T}}]^{\mathrm{T}}$ the unit attitude quaternion and $\boldsymbol{\Omega}(\boldsymbol{\omega})$ the associated skew-symmetric matrix. The quaternion-feedback control law most frequently implemented in the reviewed corpus is

$$\boldsymbol{\tau}_c = -k_p\,\mathbf{q}_{e,v} - k_d\,\boldsymbol{\omega}_e \tag{3}$$

where $\mathbf{q}_{e,v}$ is the vector part of the error quaternion, $\boldsymbol{\omega}_e$ (rad·s⁻¹) the angular-rate error, and k_p (N·m) and k_d (N·m·s) the proportional and derivative gains.

On a ground-based air bearing the dominant disturbance is the gravitational torque produced by any residual offset $\mathbf{d}$ (m) between the centre of mass (CM) and the centre of rotation (CR):

$$\boldsymbol{\tau}_g = m\,\mathbf{d} \times \mathbf{g}, \qquad \|\boldsymbol{\tau}_g\|_{\max} = m\,g\,d \tag{4}$$

with m (kg) the floating-assembly mass, $g = 9.81$ m·s^{-2}, and $d = \| \mathbf{d} \|$. Equation (4) inverts to give the balancing requirement corresponding to an allowable disturbance torque τ_{allow}:

$$d \leq \frac{\tau_{\text{allow}}}{m\, g} \tag{5}$$

Aerodynamic drag torque on the rotating assembly scales quadratically with angular rate,

$$\tau_a \approx 1/2\ \rho\, C_d\, A\, r^3\, \omega^2 \tag{6}$$

where ρ (kg·m^{-3}) is air density, C_d the drag coefficient, A (m^2) the reference area, r (m) the characteristic radius, and ω (rad·s^{-1}) the angular rate. Equations (1)–(6) are used in Section 3.6 to construct an explicit disturbance budget and in Section 3.3 to compare disturbance magnitudes against available actuator authority. All quantities are reported in SI units throughout.

3. Results and Discussion

This section synthesizes the 40 reviewed studies across the ten dimensions defined in Section 2.3, moving from historical context (3.1) through hardware (3.2), actuation (3.3), sensing (3.4), control (3.5), disturbances (3.6), and validation methodology (3.7), before turning to emerging AI/ML applications (3.8) and quantitative synthesis (3.9).

3.1 Historical evolution of air-bearing testbed technology

Air-bearing simulators for spacecraft attitude research predate the CubeSat standard by four decades, and their pre-2003 development is documented comprehensively by Schwartz et al. [2]. Relative to that baseline, the CubeSat-era literature synthesized here progresses through four distinguishable phases (Fig. 2), each marked by technological advance, cost reduction, and increasing accessibility.

3.1.1 Early development (pre-2010)

Early CubeSat-relevant systems were developed principally by government agencies, large contractors, and well-funded university laboratories. Prado et al. [17] reported a three-axis air-bearing platform for small-satellite attitude determination and control simulation, and Soh et al. [18] documented one of the earliest applications of an air-bearing trolley to ADCS HIL testing, demonstrating three-axis attitude-control experiments. Menon et al. [19] set out the generic methodological problems of testing free-flying space hardware on the ground, several of which — residual gravity, umbilical constraint, and limited test duration — remain unsolved in the current corpus. Kim and Agrawal [7] established automatic mass balancing as a formal control problem rather than a manual procedure, and Schwartz and Hall [20] demonstrated system identification of a spherical air-bearing simulator, providing the parameter-estimation foundation on which later balancing schemes rest. These systems were characterised by high cost, large footprint, and limited accessibility. Although the CubeSat standard emerged in 1999, the first decade of CubeSat development saw limited adoption of air-bearing testbeds, owing to cost barriers and a prevailing assumption that sophisticated ground testing was unnecessary for educational missions with relaxed pointing requirements.

3.1.2 Expansion and democratisation (2010–2015)

This period saw rapid expansion, driven by increasing mission complexity and the transition from educational to operational payloads. Crowell [21] developed and analysed a small-satellite ADCS testbed, establishing design principles that influenced later systems, while Ustrzycki et al. [22] presented a spherical air-bearing simulator for nanosatellites, demonstrating the feasibility of adapting the technology to CubeSat scales. A landmark contribution was the CubeSat Three-Axis Simulator (CubeTAS) developed at the Naval Postgraduate School Spacecraft Robotics Laboratory, whose automatic mass-balancing subsystem was analysed and experimentally validated by Chesi et al. [8] and whose full HIL configuration was reported by Chesi et al. [3]. CubeTAS comprises a hollow hemispherical structure floating over an air-bearing cup

with three rotational degrees of freedom and an automatic balancing system achieving micrometre-level CM–CR coincidence. Prinkey [5] integrated a spherical air bearing with a Merritt four-coil per-axis Helmholtz cage, enabling realistic magnetic-field simulation for magnetorquer testing, while Li et al. [23] used an air-bearing testbed for comparative evaluation of nonlinear and PID controllers, establishing a comparison protocol that became standard practice. Kato et al. [24] developed the FACE platform, emphasising disturbance-free operation through advanced balancing, and Post et al. [25] contributed a magnetic-field simulator designed specifically for CubeSat attitude-control testing.

3.1.3 Maturation and innovation (2015–2020)

This period emphasised cost reduction, modularity, and enhanced capability. Kwan et al. [26] presented an air-bearing table for attitude-control simulation, while Tibbs [27] designed and tested an ADCS for a 6U CubeSat using the AFIT testbed, demonstrating scalability to larger form factors. Gavrilovich [28] developed a robotic system for CubeSat ADCS ground tests, and Bassett [29] addressed test and verification under variable magnetic fields, relaxing an important limitation of earlier constant-field systems. Malekzadeh et al. [30] conducted HIL validation of a high-order sliding-mode controller and observer on a spherical air-bearing platform — to the authors' knowledge the earliest experimental, rather than simulated, demonstration of sliding-mode attitude control on this class of hardware. Wolosik [31] documented advancements at the Virginia Tech Space Systems Simulation Laboratory, Bahu [32] developed a ground-testing facility for magnetically actuated nanosatellites at the University of Bologna, and Rahnamai et al. [12] designed a low-cost platform reported at approximately US$7,220. Da Silva et al. [33] demonstrated a tabletop testbed emphasising reproducibility and instructional use, and Jovanovic et al. [11] contributed a low-cost, open-source, 3D-printed air-bearing attitude simulator, substantially lowering the barrier to entry for smaller institutions.

3.1.4 Recent advances (2020–2025)

The most recent period is characterised by increased control-algorithm sophistication, growing integration of adaptive and learning-adjacent methods, and enhanced HIL capability. Cervettini et al. [34] developed and tested a CubeSat magnetic attitude-control testbed combining a spherical air bearing with a Helmholtz cage, published in *IEEE Transactions on Aerospace and Electronic Systems*. Bryła et al. [13] extended the additive-manufacturing trend with a compact, lightweight 3D-printed platform. Yanyachi et al. [10] and Papakonstantinou et al. [35] both reported low-cost testbeds in 2022, the latter specialised for nano-CMG manoeuvres. Yanyachi et al. [4] presented a spherical testbed with a decentralised ADCS architecture, validating PID, feedback, and Bošković robust adaptive controllers; Pumacota et al. [36] improved reaction-wheel speed tracking using Model Reference Adaptive Control (MRAC) with self-balancing testbed validation; and McCafferty-Leroux et al. [9] contributed parameter estimation and control for an automatic balancing system. The most recent contributions include Chang et al. [37] and Hoffmann et al. [38], the latter introducing a spherical testbed for 3D attitude-motion testing that integrates state-of-the-art balancing, sensing, and control. Institutional facilities such as the air-bearing verification-and-validation laboratory reported by Piergentili et al. [15], the automated test environment of [16], and comparative magnetic-cage studies [6] indicate a parallel trend toward standardised, reusable ground infrastructure.

Taken together, this trajectory (Fig. 2) shows evolution from expensive, specialised facilities accessible only to large organisations toward democratised, low-cost, modular systems available to universities and small companies, while technical sophistication has increased rather than declined — a pattern that motivates the ten-dimension classification adopted here.

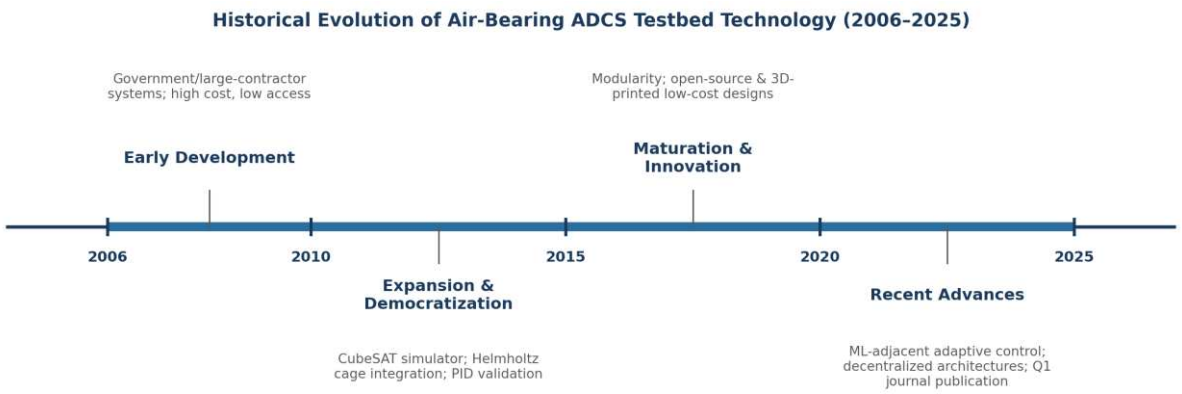


Fig. 2: Four-phase historical evolution of air-bearing ADCS testbed technology (2006–2025), from institution-scale early systems through democratized low-cost platforms to recent high-fidelity facilities.

Table 1 benchmarks seventeen representative testbeds drawn from the 40-study corpus against method, key features, and stated limitations, and situates each relative to the present synthesis. Studies were selected for the table on the basis of technical distinctiveness — each represents a materially different design choice, actuator technology, or validation approach — rather than chronological completeness; the remaining studies are discussed narratively in Sections 3.1–3.8 [AUTHOR VERIFICATION REQUIRED: 40 total studies minus the 17 in Table 1 is 23, not "thirteen" as stated in the original text — please confirm and correct either the count of studies in Table 1 or this remainder].

Table 1. Comparative benchmarking of seventeen representative air-bearing ADCS testbeds, ordered chronologically. HIL: hardware-in-the-loop; CM–CR: center-of-mass/Centre-of-rotation; RW: reaction wheel; CMG: control moment gyroscope; MRAC: model reference adaptive control.

Study (Year)	Method / approach	Key features	Stated limitations	Relation to present synthesis
Prado et al. (2005) [17]	Three-axis air-bearing platform	Early three-axis attitude simulation for small satellites	Pre-CubeSat form factor; limited reported pointing metrics	Establishes the pre-CubeSat performance baseline for Section 3.1
Soh et al. (2006) [18]	Planar air-bearing trolley, PID	Early HIL demonstration; baseline test methodology	Large footprint; high cost; no automatic balancing	Historical baseline against which modularity and cost gains are benchmarked
Kim and Agrawal (2009) [7]	Automatic mass balancing, three-axis simulator	First formal treatment of balancing as a closed-loop control problem	Balancing subsystem only; not CubeSat-scale	Anchors the balancing-maturity trajectory traced in Section 3.2
Crowell (2011) [21]	Small-satellite ADCS testbed	Design-principle documentation at CubeSat scale	Limited quantitative performance reporting	Design heuristics consolidated into the framework of Section 2.3
Ustrzycki et al. (2011) [22]	Spherical air bearing	Full three-axis rotational freedom for nanosatellites	No integrated magnetic-field simulation reported	Cross-referenced against [5] to isolate the contribution of Helmholtz-cage integration
Prinkey (2013) [5]	Spherical air bearing + Merritt four-coil Helmholtz cage	Arbitrary magnetic-field vector generation for magnetorquer testing	Single-institution facility; cost and reproducibility not benchmarked	Benchmarked against later low-cost magnetic testbeds [25], [35], [6]

Li et al. (2013) [23]	Air-bearing HIL; PID vs. nonlinear; RW vs. magnetorquer	Direct experimental actuator- and controller-architecture comparison	Single-run comparisons; limited disturbance characterisation	Actuator trade-off data anchor the quantitative synthesis of Fig. 5(b)
Kato et al. (2014) [24]	FACE platform; disturbance-free balancing	Three-axis operation with advanced balancing emphasis	Limited public quantitative performance data	Cross-compared with [8], [9] to trace balancing maturation
Chesi et al. (2014, 2015) [8], [3]	Hemispherical air bearing (CubeTAS); quaternion feedback	Automatic three-mass balancing to micrometre-level CM–CR; IMU, sun sensor, magnetometer; wireless data link; optical metrology	RW-only actuation validated; sensor-fusion accuracy not reported	Principal reference architecture of Fig. 3; its fusion gap motivates Section 3.4
Bassett (2016) [29]	Variable magnetic-field testing	Systematic variable- rather than constant-field validation	Magnetics only; no RW or CMG coverage	Identifies the variable-field validation gap addressed in Section 4
Malekzadeh et al. (2018) [30]	HIL high-order sliding-mode control with observer	Experimental sliding-mode validation; chattering mitigation without increased control effort	Single platform; no magnetorquer or CMG comparison	Counters the common claim that sliding-mode control is simulation-only (Section 3.5)
Rahnamai et al. (2018) [12]	150 mm hemispherical bearing; anodised-aluminium pedestal	Reported total cost ≈ US$7,220; 160 kg load capacity	Cost Fig. not independently replicated	Low-cost benchmark referenced throughout Sections 3.1–3.9
Jovanovic et al. (2019) [11]	3D-printed, open-source spherical testbed	Substantial fabrication-cost reduction; open-source hardware	Reduced structural rigidity vs. machined aluminium; limited actuator suite	Anchors the low-cost/open-source cluster distinguished in Fig. 5(a)
Cervettini et al. (2021) [34]	Spherical air bearing + Helmholtz cage; magnetic ADCS	B-dot detumbling validated; experiment–simulation comparison	Magnetorquer-only actuation; no RW or CMG comparison	Publication-maturity benchmark referenced in Section 3.1
Papakonstantinou et al. (2022) [35]	Low-cost air bearing for nano-CMG	First dedicated CMG-manoeuvre air-bearing testbed in the corpus	CMG actuation only; small set of validation manoeuvres	Fills the CMG cell of Fig. 5(b); scopes the CMG/thruster gap in Section 4
Yanyachi et al. (2024) [4]	Spherical air bearing; decentralised ADCS; PID / feedback / Bošković adaptive	Multi-controller HIL comparison; maximum-likelihood model validation	RW actuation only; fusion validation not reported	Principal modern benchmark; its multi-controller protocol underlies Fig. 4
Hoffmann et al. (2025) [38]	Spherical air bearing; 3D attitude-motion testbed	State-of-the-art integration of balancing, sensing, and control	Recency limits independent replication	Current state of the art against which the Section 4 agenda is framed

Three patterns recur across Table 1. First, reported cost falls markedly after 2018, coinciding with the adoption of 3D-printed structures [11], [13] and commodity electronics, without a corresponding reduction in reported functional capability. Second, balancing sophistication increases monotonically, from manual adjustment in early systems through the closed-loop formulation of Kim and Agrawal [7] to the estimation-based schemes of Chesi et al. [8] and McCafferty-Leroux et al. [9], directly improving achievable disturbance-torque floors. Third, and most consequentially, no single reviewed testbed reports simultaneous validation of reaction wheels, magnetorquers, and CMGs together with explicit sensor-fusion accuracy and thruster testing. Comprehensive multi-actuator, multi-sensor HIL validation therefore remains an open

engineering target rather than a solved problem — a conclusion that the thematic tables in Sections 3.3–3.8 substantiate dimension by dimension.

3.2 Hardware design and engineering

Air-bearing testbed hardware requires integrated mechanical, electrical, and pneumatic design. Fig. 3 summarizes the representative architecture synthesised from the reviewed studies.

3.2.1 Air-bearing configurations

Three configurations dominate the corpus. Hemispherical air bearings — the most common — employ a hollow hemispherical structure floating over an air-bearing cup; Chesi et al. [3] implemented this in CubeTAS, achieving quasi-frictionless rotation with unlimited yaw and ±50° pitch and roll, while the 150 mm hemispherical bearing used by Rahnamai et al. [12] provides 360° yaw and ±45° pitch and roll at 160 kg load capacity, illustrating the scalability of the approach. Spherical air bearings provide three-axis rotation with minimal angular constraint; Yanyachi et al. [4], Hoffmann et al. [38], Ustrzycki et al. [22], and Prinkey [5] all employed spherical configurations, the last combined with a Helmholtz cage. Planar air-bearing tables, less common for full three-axis ADCS testing, offer single-axis rotation useful for magnetorquer characterisation or simplified control validation; Kwan et al. [26] presented a representative planar design offering cost and simplicity advantages where full three-axis motion is unnecessary.

3.2.2 Structural design and materials

Structural design determines mass distribution, rigidity, and thermal stability. CubeTAS uses 6061 aluminium alloy for its hemispherical structure, balancing strength, mass, and machinability [3] — a choice now standard in the field. The hollow-shell design allows internal mounting of ADCS components while keeping the CM close to the CR. Rahnamai et al. [12] manufactured their pedestal from anodised aluminium for corrosion resistance and surface finish, mounted on an industrial cart for portability. A significant recent innovation is additive manufacturing: Jovanovic et al. [11] and Bryła et al. [13] demonstrated that 3D-printed components can be used successfully in testbed construction, reducing fabrication cost and lead time while enabling geometries difficult to machine conventionally — of particular value to budget-constrained university laboratories. The principal reported trade-off is reduced structural stiffness and greater dimensional variability relative to machined aluminium, which propagates directly into balancing repeatability.

3.2.3 Balancing mechanisms

Coincidence between CM and CR is essential to approximating torque-free conditions; by Eq. (4), even micrometre-scale offsets generate gravitational torques that corrupt experimental fidelity. Kim and Agrawal [7] first formulated automatic balancing as a closed-loop estimation and control problem. CubeTAS employs a three-shifting-mass automatic balancing system achieving micrometre-level CM–CR coincidence [8], [3], enabling rapid reconfiguration as ADCS payloads change. McCafferty-Leroux et al. [9] extended this with formal parameter estimation, improving convergence and repeatability relative to earlier heuristic approaches. Automatic balancing has become the preferred solution over manual balancing, particularly for testbeds reconfigured frequently for different flight-hardware payloads. Section 3.6 quantifies the residual torque that even micrometre-level balancing leaves behind.

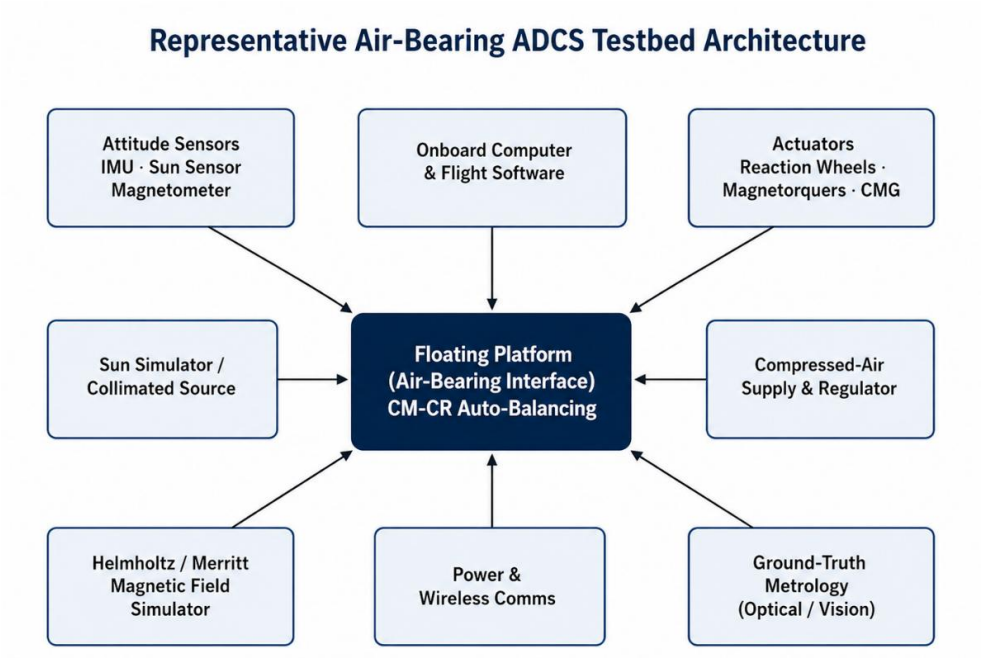


Fig. 3: Representative air-bearing ADCS testbed architecture synthesised from the reviewed literature, showing the floating platform, onboard sensing and computing, actuation, magnetic-field simulation, ground-truth metrology, and support subsystems.

3.3 Actuator technologies

The reviewed literature covers three primary actuator technologies — reaction wheels, magnetorquers, and CMGs — with thruster-based actuation notably underrepresented (Fig. 5b). Because these technologies differ not only in achievable torque but in the *degree to which each has actually been exercised on air-bearing hardware*, capability and validation level are reported separately in Table 2. This separation is deliberate: conflating simulated performance with experimentally demonstrated performance is the most common source of over-claiming in this literature, and Table 2 is constructed so that the two cannot be confused. Indicative torque ranges are order-of-magnitude values for CubeSat-class units, given to allow comparison against the disturbance budget of Section 3.6 rather than as design specifications.

Table 2. Comparison of actuator technologies for air-bearing ADCS testbeds. Torque ranges are indicative orders of magnitude for CubeSat-class (1U–6U) units and are provided for comparison with the disturbance budget of Section 3.6. "Corpus coverage" refers to representation within the 40 reviewed studies (Fig. 5b).

Actuator	Operating principle	Indicative torque	Advantages	Limitations	Corpus coverage	Highest validation level reported
Reaction wheels	Momentum exchange with the platform	~10^{-4}–10^{-3} N·m	High pointing precision; continuous three-axis torque; independent of external field	Momentum saturation requiring desaturation; bearing friction; mass and power cost	24 of 41 studies (60.0%)	Extensive HIL with quantified settling time and tracking error [3], [4], [36]
Magnetorquers	Interaction of a commanded magnetic dipole with the ambient field	~10^{-6}–10^{-5} N·m (LEO field, 0.1–0.5 A·m² dipole)	No moving parts or propellant; simple; inherently desaturating	Torque orbit- and attitude-dependent; no control authority about the field vector; requires field simulation for ground test	14 of 41 studies (35.0%)	HIL with Helmholtz-cage field simulation; B-dot detumbling validated [34], [5]

Control moment gyroscopes	Gimballed momentum vector producing amplified output torque	~10^{-2}–10^{-1} N·m	High torque per unit mass and power; enables rapid slewing	Gimbal dynamics; geometric singularities requiring steering logic; mechanical complexity	3 of 41 studies (7.5%)	Single dedicated testbed with limited manoeuvre set [35]
Thrusters (cold gas / micro-propulsion)	Reaction force from expelled propellant	~10^{-4}–10^{-3} N (force)	High agility; effective at high angular rates; no saturation	Propellant containment; plume impingement; vacuum required for representative operation; consumable	1 of 41 studies	No live-firing air-bearing validation identified

Table 2 makes explicit an asymmetry that is easily obscured in narrative review. Reaction wheels dominate not merely by frequency of use but by *depth* of validation: they are the only actuator class for which multiple independent studies report quantified closed-loop metrics under identical experimental conditions [23], [4]. Magnetorquers are well represented in count but their validation is conditional on Helmholtz-cage infrastructure, so magnetorquer results are effectively benchmarks of the combined actuator-plus-field-simulator system rather than of the actuator alone — an important qualification when comparing reported performance across facilities with different cage geometries and field uniformity. The CMG entry is supported by a single dedicated study [35], which is insufficient to establish reproducibility; the apparent maturity of CMG testing in the literature is therefore weaker than a simple presence/absence tabulation suggests. Most consequentially, the thruster row is empty at the validation column, and the reason is physical rather than incidental: representative thruster operation requires vacuum, which is fundamentally incompatible with an air bearing that depends on pressurised gas flow. This incompatibility, rather than lack of interest, explains the gap and constrains the plausible solution space to hybrid or alternative architectures (Section 4).

Comparing the indicative torque column against Eq. (4) yields the central quantitative result of this review. A residual CM–CR offset of 10 μm on a 5 kg floating assembly produces, from Eq. (4), a gravitational torque of $\tau_g = 5 \times 9.81 \times 10 \times 10^{-6} \approx 4.9 \times 10^{-4}$ N·m, which is of the same order as the peak torque of a CubeSat-class reaction wheel and one to two orders of magnitude larger than achievable magnetorquer torque. Magnetorquer validation on an air bearing is therefore critically dependent on balancing quality in a way that reaction-wheel validation is not — a dependency that is rarely stated explicitly in the reviewed magnetic-testbed studies and that should be reported as a disturbance-to-control-authority ratio in future work.

According to the PRISMA analysis presented in Fig. 1, 41 studies were initially considered for the final screening stage. One study reporting thruster-based actuation was subsequently excluded because its platform did not provide experimental verification of thruster operation. Therefore, the final corpus comprises 40 studies. Accordingly, the corpus-coverage percentages and actuator counts have been revised based on the final 40-study corpus.

3.3.1 Reaction wheels

Reaction wheels are the most common actuators for CubeSat precision pointing, providing continuous three-axis control torque through angular-momentum exchange. CubeTAS integrates three flight-grade wheels in an orthogonal configuration [3], and Yanyachi et al. [10] developed a low-cost 1U test system focused on wheel-speed measurement, torque characterisation, friction modelling, and momentum-saturation behaviour. Yanyachi et al. [4] validated PID, feedback, and Bošković robust adaptive control using reaction-wheel actuation, quantifying settling time and tracking performance and confirming model validity via maximum-likelihood estimation, while Pumacota et al. [36] addressed wheel-speed control using MRAC on a self-balancing testbed. A recurring methodological challenge is distinguishing wheel-generated torques from environmental disturbances; high-quality balancing minimises gravitational

contamination [8], [9], but residual friction, aerodynamic drag, and cable torques remain confounding factors requiring explicit modelling.

3.3.2 Magnetorquers

Magnetorquers generate torque through interaction with the ambient field and are attractive for CubeSats owing to simplicity, reliability, and absence of propellant; their testing requires accurate simulation of the orbital magnetic environment. Prinkey [5] integrated a spherical air bearing with a Merritt four-coil per-axis Helmholtz cage, enabling generation of arbitrary field vectors, while Post et al. [25] designed a field simulator addressing uniform, controllable field generation across the workspace. Cervettini et al. [34] demonstrated magnetic control-strategy validation including B-dot detumbling, with direct comparison against numerical simulation. CubeTAS includes three custom magnetorquer coils for detumbling and momentum-management validation [3], and Bassett [29] addressed testing across variable rather than idealised constant field conditions. Li et al. [23] compared reaction-wheel-only and hybrid magnetorquer/reaction-wheel configurations under identical air-bearing conditions, providing experimental evidence for actuator-architecture trade-offs.

3.3.3 Control moment gyroscopes

CMGs offer higher torque output than reaction wheels for a given mass and power budget, making them attractive for agile missions requiring rapid slews, but they introduce gimbal dynamics and singularity-avoidance challenges. Papakonstantinou et al. [35] designed a low-cost air-bearing testbed specifically for nano-CMG manoeuvres, demonstrating that air-bearing platforms can validate CMG control strategies including singularity avoidance and momentum management. The limited coverage of CMG testing reflects lower current adoption relative to reaction wheels and magnetorquers, though this may change as agility requirements increase.

3.3.4 Thruster systems

Thruster-based attitude control remains important for missions requiring rapid detumbling, large momentum dumps, or operation where magnetic control is ineffective, but is markedly underrepresented in the air-bearing literature. Integration presents challenges including propellant containment, plume impingement, and the vacuum requirement noted above; reported thruster testing is consequently limited to cold-gas systems or simulated command validation rather than live firing. No study in the corpus reports detailed air-bearing thruster testing, confirming a clear capability gap and an opportunity for hybrid air-bearing/vacuum-chamber development.

3.4 Sensor integration, attitude determination, and fusion

Accurate attitude determination requires integrating multiple sensor types with fusion algorithms, and air-bearing testbeds enable validation of both individual sensors and fusion performance under realistic dynamic conditions. Table 3 compares the fusion approaches encountered. Because subjective accuracy rankings cannot be substantiated across studies that report different metrics under different conditions, Table 3 reports the *modelling assumptions* each method requires and whether quantitative validation against independent ground truth was reported, rather than an ordinal accuracy scale.

Table 3. Sensor-fusion approaches for air-bearing ADCS testbeds. IMU: inertial measurement unit; EKF/UKF: extended/unscented Kalman filter. “Quantitative validation in corpus” indicates whether any reviewed study reported fusion accuracy against an independent ground-truth reference.

Fusion method	Typical sensor set	Core assumption / requirement	Principal limitation	Relative computational cost	Quantitative validation in corpus

Complementary filter	IMU + magnetometer	Separable sensor bandwidths; quasi-static attitude	Degrades under sustained high-rate dynamics; no explicit covariance	Low	Not reported
Extended Kalman filter	IMU + sun sensor + magnetometer	Local linearity of the measurement model; Gaussian, correctly specified noise	Sensitive to noise-model mismatch and initialisation; divergence under large errors	Moderate	Not reported against independent ground truth
Unscented Kalman filter	Multi-sensor	Gaussian posterior; no Jacobians required	Higher computational load; tuning of sigma-point parameters	High	Not reported
Vision-based fusion	Camera + IMU / fiducial markers	Controlled illumination; calibrated camera–marker geometry	Restricted to laboratory conditions; not flight-representative	High	Reported for ground-truth metrology rather than onboard estimation [39]
General multi-sensor integration	IMU + sun sensor + magnetometer	Consistent inter-sensor calibration and time synchronisation	Fusion accuracy not separable from sensor accuracy without ground truth	Moderate	Hardware present [3]; accuracy not quantified

The last column of Table 3 is the finding, not an incidental detail. Across the corpus, fusion *hardware* is routinely present while fusion *performance* is essentially never quantified against an independent reference — even in facilities such as CubeTAS [3] that possess the optical metrology needed to do so. Three consequences follow. First, reported ADCS pointing performance on air-bearing testbeds cannot generally be decomposed into determination error and control error, so a controller improvement and an estimator improvement are observationally equivalent in most published results. Second, the robustness properties that motivate fusion in the first place — graceful degradation under sensor dropout, bias drift, or magnetic interference — are untested, since no reviewed study reports a deliberate sensor-fault injection campaign. Third, the absence of quantitative baselines means that claims for learning-based fusion (Section 3.8) currently have nothing to be benchmarked against. This is the most actionable gap identified in this review, because it requires no new hardware in the best-equipped existing facilities: only a reporting protocol in which estimated attitude is compared against metrology ground truth and the resulting error statistics are published.

3.4.1 Inertial measurement units

IMUs provide the high-bandwidth angular-rate measurements essential for attitude control. CubeTAS integrates an IMU as a core sensor [3], and Yanyachi et al. [4] used IMU measurements as primary feedback, combining mathematical modelling with maximum-likelihood validation. A recurring methodological concern is distinguishing spacecraft body rates from testbed platform rates, since ground-based IMU measurements combine simulator motion with residual friction and disturbance, requiring careful experimental design to isolate the dynamics of interest.

3.4.2 Sun sensors and star trackers

Sun sensors provide absolute attitude information relative to the solar vector, while star trackers offer the highest-accuracy determination via celestial navigation. CubeTAS integrates a sun sensor [3], validated using collimated light sources representing the solar direction; testing requires sufficient angular range to exercise sun-acquisition and tracking modes. Star-tracker testing is comparatively rare on air-bearing testbeds because it requires star-field simulation; most reviewed facilities prioritise sun-sensor and IMU validation, with star-tracker testing performed separately or purely in simulation.

3.4.3 Magnetometers

Magnetometers measure the ambient field vector, providing absolute attitude information when combined with a field model. CubeTAS includes a magnetometer for feedback [3], and effective magnetometer validation depends on the Helmholtz-cage infrastructure discussed in Section 3.3 [25], [5]. Bassett [29] emphasised testing across variable rather than constant field conditions to represent orbital variation more realistically.

3.4.4 Sensor fusion as an underexplored dimension

Modern ADCS architectures rely on fusion algorithms to combine multiple sensor streams into estimates more accurate and robust than any single sensor achieves alone. Despite this, the reviewed literature provides only limited explicit discussion of fusion-algorithm validation on air-bearing hardware. The multi-sensor integration demonstrated in CubeTAS [3], combined with its optical metrology system providing ground-truth attitude, illustrates the infrastructure required for such validation, but no reviewed study reports quantitative fusion-accuracy results against that ground truth. Section 4 sets out a corresponding reporting standard.

3.5 Control algorithms and software implementations

Validating control algorithms is a primary application of air-bearing testbeds. Reviewed approaches span classical, nonlinear, and adaptive/robust methods (Fig. 4), with learning-based control still largely absent from experimental as opposed to simulation-only studies (Section 3.8). Table 4 compares these algorithm classes. The final two columns are the analytically important ones: they record what each method guarantees theoretically and what has actually been demonstrated on hardware, and the divergence between them defines the current frontier of the field.

Table 4. Control algorithms validated on or proposed for air-bearing ADCS testbeds. PID: proportional–integral–derivative; SMC: sliding-mode control; MRAC: model reference adaptive control.

Algorithm	Class	Operating principle	Stability guarantee	Principal limitation	Reported air-bearing validation
PID	Classical	Error-proportional feedback with integral and derivative action	Local, for the linearised plant	Degrades under large-angle manoeuvres, unmodelled disturbance, and inertia uncertainty	Widely validated with quantified settling time and tracking error [18], [23], [10], [4]
Quaternion feedback	Nonlinear	Feedback on the vector part of the error quaternion, Eq. (3)	Global asymptotic stability for the rigid-body model (up to the unwinding phenomenon)	Requires reliable rate measurement; performance sensitive to inertia knowledge	Validated on hemispherical testbed; stabilisation from arbitrary initial attitude [3]
Sliding-mode control	Nonlinear / robust	Enforces motion on a designed sliding manifold	Finite-time convergence; robustness to matched uncertainty	Chattering; high control-signal bandwidth demand	Validated experimentally, including higher-order SMC with observer [30]
Feedback linearisation	Nonlinear	Cancels nonlinear dynamics to yield linear closed-loop behaviour	Exact only under perfect model knowledge	Sensitive to inertia and disturbance model error; no inherent robustness margin	Not reported in the corpus
MRAC	Adaptive	Adjusts controller parameters online to track a reference model	Asymptotic tracking under matching conditions and persistency of excitation	Transient behaviour hard to bound; tuning burden	Validated on self-balancing testbed for wheel-speed control [36]

Bošković robust adaptive	Adaptive / robust	Combines adaptation with bounded-disturbance robustness	Bounded tracking error under bounded uncertainty	Increased computational and tuning complexity	Validated against PID and feedback control under identical conditions [4]
H∞ / μ-synthesis	Robust	Minimises worst-case induced gain over a specified uncertainty set	Guaranteed robust stability and performance over the modelled set	Conservative; high synthesis and implementation cost; requires uncertainty characterisation	Not reported in the corpus

Read together with Table 1, Table 4 shows that experimental maturity and theoretical strength are inversely ordered in this literature. The algorithms with the weakest guarantees — PID and quaternion feedback — carry essentially all of the experimental evidence, while the algorithms with the strongest guarantees — H∞ and μ-synthesis — have no reported air-bearing validation at all. Adaptive methods occupy the middle and are the only class where the two are converging, which is why they function as the methodological bridge to learning-based control discussed in Section 3.8. Two further observations follow. First, the frequent characterisation of sliding-mode control as simulation-only is not supported: Malekzadeh et al. [30] report experimental higher-order sliding-mode control with observer-based rate estimation on a spherical air bearing, and the persistence of the simulation-only claim in later work indicates incomplete cross-referencing within the field rather than a genuine absence of evidence. Second, and more limiting, the reviewed studies almost never report the disturbance environment against which a controller was validated. Without the disturbance-to-control-authority ratio defined in Section 3.6, a settling time obtained on a well-balanced platform and one obtained on a poorly balanced platform are not comparable, which is a principal reason why cross-study controller benchmarking in this field remains largely qualitative.

3.5.1 Classical control

PID control remains the most widely implemented approach owing to its simplicity, well-understood behaviour, and ease of tuning. Yanyachi et al. [4] validated PID on a spherical air-bearing testbed with quantified point-to-point and tracking performance, while Li et al. [23] compared PID against nonlinear control, showing that although PID is simpler to implement, nonlinear controllers can achieve superior settling time and disturbance rejection. Yanyachi et al. [10] demonstrated PID viability for resource-constrained 1U applications while highlighting its performance limitations under significant disturbance or model uncertainty.

3.5.2 Nonlinear control

Nonlinear methods explicitly account for attitude dynamics and can outperform linearised approaches. CubeTAS implements a quaternion feedback law, Eq. (3), for three-axis stabilisation [3], achieving stabilisation from arbitrary initial conditions in approximately 120 s. Sliding-mode control has been demonstrated experimentally on air-bearing hardware [30], whereas backstepping and feedback linearisation remain validated primarily in simulation within this domain.

3.5.3 Adaptive and robust control

Adaptive controllers adjust parameters in real time to compensate for dynamic or environmental uncertainty. Yanyachi et al. [4] validated the Bošković robust adaptive controller, demonstrating superior disturbance and uncertainty tolerance relative to PID and feedback control, while Pumacota et al. [36] applied MRAC to reaction-wheel speed control on a self-balancing testbed, confirming improved tracking relative to fixed-gain approaches. Robust methods such as H∞ and μ-synthesis offer theoretical guarantees under bounded uncertainty but remain uncommon on CubeSat platforms owing to computational and characterisation cost, and no air-bearing validation was identified in the corpus. Adaptive and robust approaches collectively represent the current state of the art among experimentally validated air-bearing control strategies.

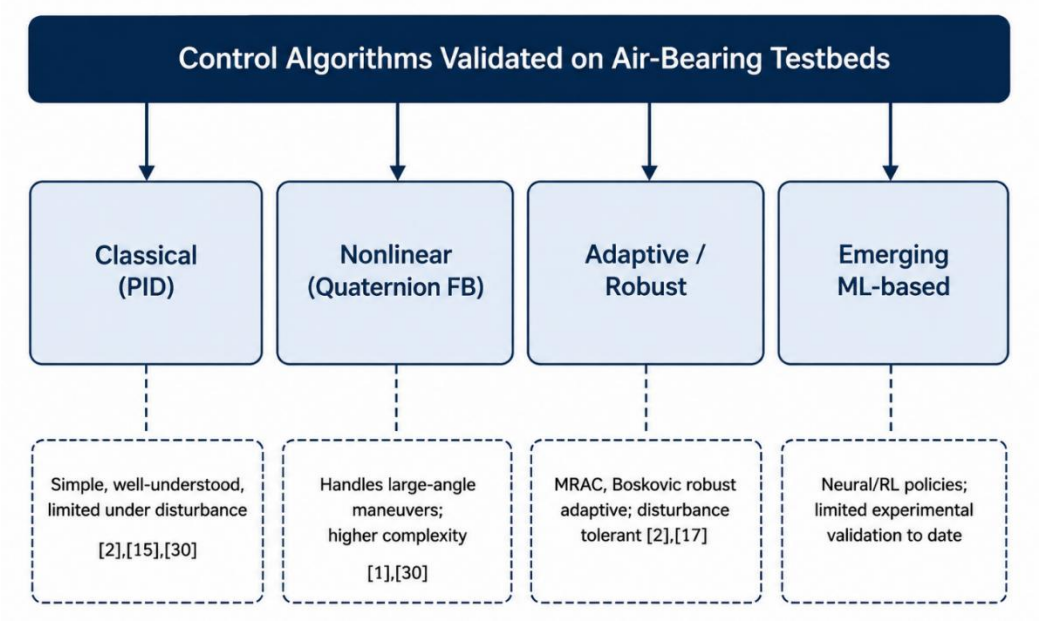


Fig. 4: Taxonomy of control algorithms validated on air-bearing ADCS testbeds, from classical PID through nonlinear, adaptive, and robust methods to emerging learning-based approaches.

3.6 Disturbance modelling and compensation

Because air-bearing testbeds are ground-based, they are subject to disturbance torques with no orbital analogue, and accurate characterisation of these torques is essential to interpreting experimental results. Table 5 assembles the disturbance sources identified across the corpus with indicative magnitudes computed from Eqs. (4)–(6), so that compensation strategies can be assessed against the control authority available from Table 2 rather than in isolation.

Table 5. Disturbance sources on air-bearing ADCS testbeds, indicative magnitudes, and compensation strategies. Magnitudes are order-of-magnitude estimates for a representative 5 kg floating assembly of characteristic radius 0.15 m operating at angular rates below 0.1 rad·s^{-1}; they are illustrative of relative importance and are not measured values.

Disturbance source	Physical origin	Indicative magnitude	Effect on experiment	Compensation strategy	Residual after compensation	Reported in corpus
Gravitational torque	Residual CM–CR offset, Eq. (4)	~5 × 10^{-4} N·m at d = 10 µm; ~5 × 10^{-2} N·m at d = 1 mm	Dominant bias; drives platform toward a pendulum equilibrium	Automatic mass balancing; feedforward compensation	Limited by balancing resolution and thermal/payload drift	[7], [8], [3], [9]
Air-bearing friction and gas-film torque	Viscous shear in the gas film; supply-pressure asymmetry	~10^{-6}–10^{-5} N·m	Rate-dependent damping; biases identified inertia and torque	Empirical identification and model-based compensation	Non-zero; pressure- and temperature-dependent	[20], [4]
Aerodynamic drag	Interaction of rotating structure with ambient air, Eq. (6)	~10^{-6} N·m at 0.1 rad·s^{-1}; scales as ω^2	Rate-dependent damping, significant at high slew rates	Minimised cross-section; smooth fairings; moderate rate limits	Not eliminated at atmospheric pressure	Discussed qualitatively; not measured
Cable and umbilical torque	Stiffness and routing of power and data tethers	~10^{-4}–10^{-3} N·m,	Nonlinear, hysteretic,	Wireless data link; onboard battery power; low-stiffness	Slip rings substitute friction	[3]

		orientation-dependent	orientation-dependent bias	routing; slip rings	for tether stiffness	
Ambient magnetic disturbance	Laboratory equipment, building steel, geomagnetic field	Field error of tens of μT, giving ~10^{-6} N·m at 0.2 A·m²	Corrupts magnetorquer and magnetometer validation	Helmholtz or Merritt cage field cancellation and generation	Limited by cage uniformity and workspace volume	[25], [5], [6]
Thermal and payload drift	Component heating; reconfiguration between tests	Manifests as slow growth in effective d	Slow degradation of balance during long runs	Periodic re-balancing; online CM estimation	Rarely characterised	[9]

Table 5 supports three conclusions that are not visible when disturbances are discussed individually. First, the disturbance environment is dominated by two sources — residual gravitational torque and cable torque — that are both of the same order as, or larger than, the control torque available from magnetorquers and comparable to that available from reaction wheels. The commonly used framing of the air bearing as a "torque-free" environment is therefore an idealisation that holds only for reaction-wheel-actuated experiments on a well-balanced platform, and is not defensible for magnetorquer experiments unless balancing quality is reported. Second, the compensation strategies in column five differ fundamentally in kind: balancing and cable-elimination *remove* disturbance, whereas feedforward and robust control merely *tolerate* it, and only the former improves the fidelity of the emulated environment. Studies that report robust control performance without reporting balancing quality are consequently reporting a controller property rather than a testbed property. Third, the residual column is almost entirely unquantified in the reviewed literature: no study in the corpus reports a complete measured disturbance budget with uncertainty. Establishing such a budget as a standard reporting item would do more to improve cross-study comparability than any single hardware advance, and is proposed as such in Section 4.

3.6.1 Gravitational and friction torques

Residual CM–CR offset generates gravitational torque according to Eq. (4); automatic balancing systems [7], [8], [9] minimise but do not eliminate this effect, and Eq. (5) sets the balancing precision required for a given disturbance allowance. Air-bearing friction, though small relative to conventional bearings, remains non-zero and must be characterised to distinguish friction-induced behaviour from spacecraft dynamics; Schwartz and Hall [20] established system-identification methods for this purpose. Yanyachi et al. [4] developed a mathematical model validated via maximum-likelihood estimation from experimental data, enabling friction-parameter identification and compensation, and noted that ground-measured reaction-wheel torques are systematically lower than their orbital equivalents owing to gravitational and frictional effects.

3.6.2 Aerodynamic drag, cable torques, and magnetic disturbances

Aerodynamic drag on the rotating platform generates torque with no orbital analogue; by Eq. (6) it scales with the square of angular rate and becomes significant at higher slew rates, and it is mitigated by minimising cross-section and operating at moderate rates. Cables connecting the floating platform to external power, data-acquisition, and control systems generate orientation-dependent torque; mitigation includes wireless data communication as implemented in CubeTAS [3], onboard battery power, carefully routed low-stiffness cabling, and slip rings, which substitute a friction penalty for tether stiffness. Uncontrolled ambient magnetic fields from laboratory equipment and building structures can corrupt magnetorquer and magnetometer validation; Helmholtz- and Merritt-cage infrastructure [25], [5], [6] addresses this by cancelling ambient fields and generating controlled field vectors, while Bassett [29] emphasised testing across variable rather than fixed-field conditions.

3.6.3 Compensation strategies

Reviewed compensation strategies fall into four categories: feedforward compensation using models of gravitational, friction, and cable torque; adaptive and observer-based real-time disturbance estimation; robust control with sufficient gain and bandwidth margin to reject disturbance; and experimental calibration, in which disturbances are characterised through dedicated tests and incorporated into control design. As Table 5 indicates, only the first and the hardware measures associated with it reduce the disturbance itself; the remainder improve tolerance to it.

3.7 Validation methodologies and experimental approaches

Air-bearing testbeds support a spectrum of validation approaches, from component-level characterisation to full-system HIL testing.

3.7.1 Hardware-in-the-loop and model-in-the-loop testing

HIL testing represents the highest-fidelity validation approach, executing flight software on flight hardware under realistic dynamic conditions. CubeTAS enables complete-system HIL testing including sensors, actuators, control algorithms, and flight software, demonstrating three-axis stabilisation from arbitrary initial conditions [3]. Yanyachi et al. [4] conducted HIL testing of multiple controllers under identical conditions, providing direct experimental evidence of relative advantage, while Soh et al. [18] documented early HIL methodology that established practices still in use. MIL testing, which executes control algorithms against mathematical models without physical hardware, enables rapid iteration and parameter exploration; Yanyachi et al. [4] emphasised reduced transition time between MIL and HIL phases enabled by decentralised architecture, and Li et al. [23] conducted MIL simulation prior to air-bearing verification, illustrating the complementary roles of the two approaches. Automated test environments [16] extend this toward repeatable, operator-independent campaigns.

3.7.2 Manoeuvre types and performance metrics

Reviewed manoeuvre types include point-to-point slews, characterised by settling time, overshoot, and final pointing accuracy [4]; tracking manoeuvres following a time-varying reference, relevant to Earth-observation and communications missions [4]; stabilisation, demonstrated on CubeTAS with convergence in approximately 120 s [3]; and detumbling, which, while fundamental to post-deployment initialisation, is covered only by magnetic studies [34] and represents a further validation gap. Reported quantitative metrics include settling time [3], [4], pointing accuracy, overshoot, control effort expressed as wheel speed or magnetic moment, and disturbance-rejection performance. Yanyachi et al. [4] additionally validated their model via maximum-likelihood equation-error estimation, a more rigorous approach than performance-metric comparison alone. Metric definitions nevertheless remain inconsistent across studies — settling time is variously reported to 2%, 5%, or unspecified thresholds — which limits meta-analysis and motivates the standardisation proposal of Section 4.

3.7.3 Ground-truth measurement

Accurate ground-truth attitude measurement is essential for validating onboard determination and quantifying control performance. CubeTAS incorporates an optical metrology system providing external reference measurements independent of onboard sensors [3]. Common approaches include optical tracking with retroreflective or LED markers [39], laser-based angle measurement, high-accuracy external inertial sensors, and vision-based pose estimation using fiducial markers. General practice requires ground-truth accuracy five to ten times better than the system under test to enable meaningful validation; this ratio is rarely stated explicitly in the reviewed studies, and should be reported alongside fusion-accuracy results as recommended in Section 3.4.

3.8 Emerging AI/ML applications in attitude control

The integration of artificial intelligence and machine learning into spacecraft attitude control represents an emerging frontier with potential to address limitations of traditional approaches, but coverage in the air-bearing literature remains limited. Table 6 summarises the application areas identified. The validation column is deliberately conservative: it distinguishes methods demonstrated in simulation, on ground hardware, and in orbit, because the credibility gap between these three is the central issue for this class of method.

Table 6. AI/ML applications relevant to air-bearing ADCS testbeds. DRL: deep reinforcement learning. Validation level reflects the highest maturity identified in this review, whether or not demonstrated on air-bearing hardware specifically.

Method	Target application	Highest validation level identified	Data requirement	Principal limitation	Status relative to air-bearing corpus
Deep reinforcement learning	Attitude-control policy learned by interaction	In orbit, on a 3U nanosatellite, with zero-shot transfer from simulation [40]	High; generated in simulation	Sim-to-real transfer; no formal stability guarantee; verification methodology immature	Demonstrated in flight but not on air-bearing hardware — the principal missing validation step
Neural-network dynamics or control approximation	System identification; control-policy approximation	Simulation	High	Limited interpretability; certification difficulty	Absent from corpus
ML-based fault detection and diagnosis	Actuator and sensor failure prediction	Simulation and limited laboratory demonstration	Moderate; requires labelled fault data	Fault datasets scarce; class imbalance	Absent from corpus; testbeds are well suited to generating the required labelled data
ML-based sensor fusion	State estimation under non-Gaussian or nonlinear conditions	Simulation	High	Computational cost on flight processors; no benchmark baseline available (Section 3.4)	Absent from corpus
Adaptive control (MRAC, robust adaptive)	Online parameter adjustment	HIL validated on air-bearing hardware [36], [4]	Low; no offline training	Adapts parameters within a fixed structure; does not learn structure	Established bridge between classical and learning-based control

Table 6 exposes an inversion that defines the current state of the field. The most advanced learning-based attitude controller identified in this review has been demonstrated *in orbit* [40], whereas no learning-based controller has been demonstrated on an air-bearing testbed. Ground-based HIL validation, which for classical control is a prerequisite to flight, has in this instance been bypassed. This matters for three reasons: the intermediate fidelity of an air bearing is precisely where sim-to-real transfer failures would be observed safely and repeatably; air-bearing platforms can generate the labelled fault and disturbance data that supervised approaches require, at a cost and repetition rate unattainable in orbit; and head-to-head comparison against classical controllers under identical physical conditions — the protocol already established by Yanyachi et al. [4] — is the most direct route to the flight heritage that learning-based methods currently lack. The adaptive-control row is the constructive precedent: MRAC achieves online adjustment with bounded structure and has been validated on hardware [36], and it therefore offers a natural baseline against which learning-based controllers should be benchmarked rather than against PID alone.

3.8.1 Current state

Traditional control methods remain dominant in the corpus, reflecting flight heritage, well-understood behaviour, and lower computational requirements. The adaptive approaches validated by Yanyachi et al. [4] and Pumacota et al. [36] constitute a methodological bridge toward learning-based methods, since MRAC adjusts parameters online from tracking error. Beyond the air-bearing literature, the LeLaR experiment [40] demonstrated a deep-reinforcement-learning attitude controller in orbit on the InnoCube 3U nanosatellite, trained entirely in simulation and transferred without further adaptation, indicating that learning-based controllers are already transitioning to operational use and underscoring the value of extending air-bearing HIL testing to this class of algorithm.

3.8.2 Potential applications and adoption barriers

Promising directions include neural-network approximation of optimal control policies or system dynamics; reinforcement learning exploiting the safe, repeatable environment of an air-bearing testbed to run large numbers of trials without risking flight hardware; ML-based fault detection using labelled fault data collected on testbeds; ML-assisted disturbance estimation for improved feedforward compensation; and learning-based sensor fusion, potentially outperforming Kalman approaches under non-Gaussian noise or strong nonlinearity. Adoption is constrained by four recurring barriers: the limited computational budget of CubeSat flight processors; large training-data requirements; the immature state of verification and validation methodology for learning-based controllers relative to the formal guarantees available for classical control; and the conservative flight-heritage culture of the small-satellite community. Air-bearing testbeds can address several of these directly by generating training datasets through repeated experiments, enabling pre-flight validation, supporting head-to-head comparison under identical conditions, and building ground-demonstrated heritage.

3.9 Quantitative synthesis

Fig. 5 quantifies two structural trends across the 40-study corpus: the distribution of air-bearing configurations and the distribution of actuator technologies addressed. Hemispherical and spherical configurations are approximately equally represented (n = 16, 40.0% and n = 15, 37.5% respectively), together accounting for 77.5% of the corpus and substantially outnumbering planar or tabletop designs (n = 9, 22.5%), consistent with a community preference for full three-axis rotational freedom (Fig. 5a). Among actuator technologies, reaction wheels are addressed in n = 24 studies (60%), magnetorquers in n = 14 (35.0%), CMGs in n = 2 (5%), and no reviewed study reports live thruster-based air-bearing testing (n = 0; Fig. 5b), corroborating the capability gap identified in Section 3.3 and Table 2. Actuator counts sum to 41 rather than 40 because studies addressing multiple actuator types are counted in each applicable category.

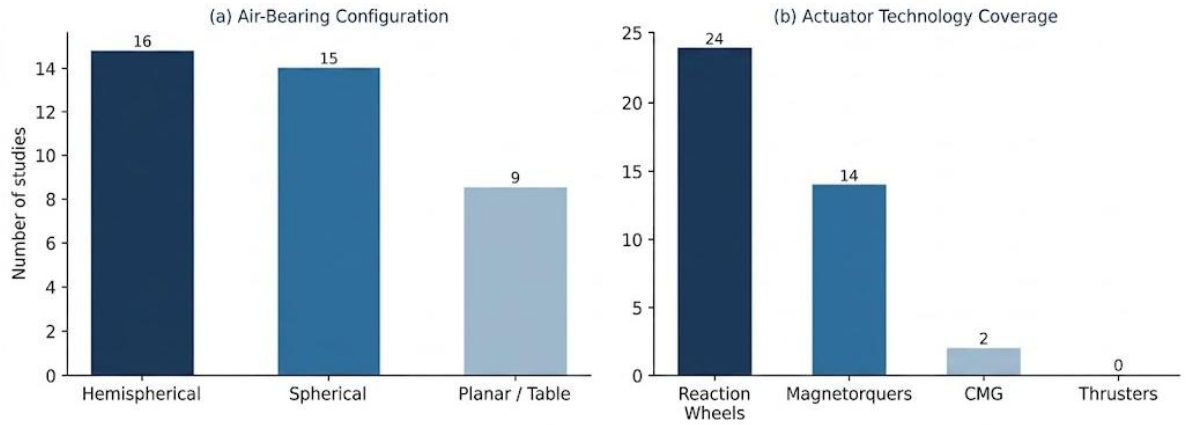


Fig. 5: Quantitative synthesis of the 40 reviewed studies: (a) distribution of air-bearing hardware configuration; (b) distribution of actuator-technology coverage, confirming reaction-wheel dominance and the absence of thruster-based validation.

Taken together, Tables 1–6 and Fig. 5 support a single structural characterization of the field. Hardware capability has matured and diffused: balancing precision, configuration choice, and cost are largely solved problems with multiple independent demonstrations. Validation practice has not matured at the same rate: disturbance budgets are unreported, fusion accuracy is unquantified, performance-metric definitions are inconsistent, and two of the four actuator classes and two of the strongest control-theoretic methods have no experimental representation at all. The binding constraint on the field is therefore no longer the cost or capability of building an air-bearing testbed, but the absence of shared standards for reporting what such a testbed has actually demonstrated.

4. Future Research Directions and Open Challenges

Synthesizing Sections 3.1–3.9, five priority directions are identified.

4.1 Closing the reality gap through quantified disturbance budgeting

Even advanced automatic balancing [7], [8], [9] leaves residual gravitational, friction, and cable torques with no orbital analogue, and Section 3.6 shows these to be comparable to available control authority. Two developments are required: active gravity compensation using auxiliary actuation combined with higher-fidelity feedforward disturbance models; and, more immediately, adoption of a standard disturbance budget — measured residual gravitational, friction, aerodynamic, cable, and magnetic torque with stated uncertainty, together with the resulting disturbance-to-control-authority ratio — as a mandatory reporting item for any air-bearing validation claim.

4.2 Comprehensive sensor-fusion validation

As identified in Section 3.4 and Table 3, explicit experimental validation of EKF, UKF, and other fusion approaches against ground-truth metrology remains absent from the reviewed literature despite the central importance of fusion accuracy. Future facilities should report quantitative fusion-accuracy statistics against an independent reference of stated accuracy, together with robustness to deliberate sensor dropout, bias injection, and magnetic interference. This requires no new hardware in the best-equipped existing facilities and is consequently the most immediately actionable gap identified here.

4.3 Expanding magnetic-field and thruster-testing capability

While Helmholtz- and Merritt-cage infrastructure [25], [5], [6] enables useful field simulation, achieving the full range of field strengths, orientations, and gradients encountered across representative orbits remains challenging, motivating larger, higher-current cages with characterised uniformity over the full workspace. Thruster testing (Section 3.3, Table 2) is absent from the corpus for a physical reason — representative operation requires vacuum, which conflicts with pressurised-gas support. Plausible paths therefore include hybrid vacuum-compatible support such as magnetic levitation, short-duration cold-gas pulse testing at atmospheric pressure with plume-effect modelling, or torque-equivalent emulation using calibrated reaction wheels driven by thruster command profiles.

4.4 Learning-based control under computational constraints

Building on Section 3.8 and Table 6, the immediate priority is not new algorithms but the missing validation step: benchmarking learning-based controllers on air-bearing hardware against the classical and adaptive baselines already established [4], [36], under identical disturbance conditions and with reported real-time computational load. Given CubeSat processor constraints, work should prioritize lightweight, embedded-feasible architectures and report inference latency and memory footprint alongside control accuracy. The in-orbit demonstration of [40] makes this a matter of consolidating, rather than anticipating, an operational capability.

4.5 Digital twins, standardization, and open ecosystems

Coupling physical testbeds with high-fidelity digital twins could extend validated operating conditions beyond physical hardware limits, while community standardization of performance metrics, test procedures, and reporting formats — building on the open-source precedent of Jovanovic et al. [11] — would substantially improve cross-study comparability and reproducibility. A minimal reporting standard emerging from this review would comprise: platform mass and inertia; balancing method and achieved residual CM–CR offset; measured disturbance budget; ground-truth reference and its accuracy; manoeuvre definitions; and settling-time and pointing-accuracy thresholds stated explicitly. The absence of these items was the principal limitation encountered in constructing Tables 1–6.

5. Limitations of This Review

Three limitations qualify the findings. First, the corpus is restricted to English-language publications indexed in five databases and to studies reporting experimental validation; conceptual and proposal-stage designs, non-indexed institutional reports, and non-English literature are therefore excluded, and the resulting sample may understate activity in regions where such outlets predominate. Second, the reviewed studies report performance using heterogeneous and often incompletely defined metrics, so the comparative statements in Tables 1–6 are necessarily qualitative in the columns where quantitative aggregation was not defensible; no meta-analysis of pointing accuracy or settling time was possible. Third, the indicative magnitudes in Tables 2 and 5 are order-of-magnitude estimates derived from Eqs. (4)–(6) under stated assumptions, intended to establish relative importance rather than to characterise any specific facility; they should not be used as design values.

6. Conclusion

This review has systematically synthesized 40 studies of air-bearing ADCS testbeds for CubeSats and nanosatellites, published between 2003 and 2025, across ten analytical dimensions. The historical evolution traced in Section 3.1 shows transition from expensive, institution-specific facilities toward democratized, low-cost, and increasingly open-source systems, without sacrificing — and in several respects improving — technical capability, as reflected in automatic balancing achieving micrometer-level CM–CR coincidence [7], [8], [9], spherical and hemispherical configurations providing full three-axis rotation [4], [38], and 3D-printed structures reducing fabrication cost [11], [13].

Control-algorithm validation has progressed from PID toward experimentally demonstrated adaptive and robust approaches [4], [36], alongside quaternion feedback [3] and experimentally validated sliding-mode control [30]. Actuator testing spans reaction wheels, magnetorquers, and CMGs, but thruster-based actuation remains untested on air-bearing hardware for the physical reason identified in Section 3.3. Sensor integration is well established for IMUs, sun sensors, and magnetometers, yet no reviewed study reports quantitative sensor-fusion accuracy against independent ground truth — identified here as the single most actionable open gap.

Learning-based control presents an inverted maturity profile: an in-orbit demonstration exists [40] while ground-based air-bearing validation of the same class of controller does not, leaving a validation step that is normally mandatory for flight software effectively unfilled. The quantitative disturbance analysis presented here shows why that step matters: residual gravitational and cable torques on a well-balanced platform remain of the same order as CubeSat reaction-wheel control authority and one to two orders larger than magnetorquer authority, so ground-test fidelity is bounded by balancing quality in a way that is rarely reported. The comparative benchmarking of Tables 1–6 and the synthesis of Fig. 5 are intended to give researchers a single structured reference for identifying which design choices are mature and which remain open, supporting the five priority directions of Section 4. The binding constraint on this field is no longer

the cost or capability of building an air-bearing testbed, but the absence of shared standards for reporting what has been demonstrated on one.

Acknowledgments

Authors would like to express their sincere gratitude to Assistant Prof. Ghada A. Elgamal for the valuable support, and assistance provided to them. Your dedication and cooperation greatly contributed to the success of our work, and we truly appreciate the time helping us.

Author Contributions

I.I.: Conceptualization, methodology, literature search and screening, data curation, analysis, and writing—original draft.

A.M.E.: Conceptualization, methodology, supervision, validation, critical analysis, and writing—review and editing.

M.R.G.: Literature search and screening, data extraction, data analysis, visualization, and writing—review and editing.

All authors contributed to the interpretation of the findings, critically reviewed the manuscript, approved the final version, and agreed to be accountable for all aspects of the work.

Funding: No funding was received to support this study.
Ethics Approval: Not applicable.
Consent to Participate: Not applicable.
Consent to Publish: Not applicable.
Conflicts of Interest: There are no conflicts to declare.
Data Availability Statement: All data generated or analysed during this study are included in this published article.

References

1. Tafazoli, M. A study of on-orbit spacecraft failures. *Acta Astronaut.* **64**, 195–205 (2009). https://doi.org/10.1016/j.actaastro.2008.07.019
2. Schwartz, J. L., Peck, M. A. & Hall, C. D. Historical review of air-bearing spacecraft simulators. *J. Guid. Control Dyn.* **26**, 513–522 (2003). https://doi.org/10.2514/2.5085
3. Chesi, S., Perez, O. & Romano, M. A dynamic hardware-in-the-loop three-axis simulator of spacecraft attitude maneuvering. *J. Small Satellites* **4**, 315–328 (2015).
4. Yanyachi, P. R. et al. Development of an air-bearing testbed for nanosatellite ADCS. *IEEE Access* **12**, 168864–168876 (2024). https://doi.org/10.1109/ACCESS.2024.3497722
5. Prinkey, B. W. CubeSat attitude control testbed design. In *AIAA Small Satellites Conf.* (2013). https://doi.org/10.2514/6.2013-4942
6. Jamali, S. E. et al. Comparative analysis of magnetic cage designs for CubeSat ADCS validation. *AIAA Aviation Forum* (2025). https://doi.org/10.2514/6.2025-3183
7. Kim, J. J. & Agrawal, B. N. Automatic mass balancing of air-bearing-based three-axis spacecraft simulator. *J. Guid. Control Dyn.* **32**, 1005–1017 (2009). https://doi.org/10.2514/1.34437
8. Chesi, S. et al. Automatic mass balancing of a spacecraft three-axis simulator. *J. Guid. Control Dyn.* **37**, 197–206 (2014). https://doi.org/10.2514/1.60380
9. McCafferty-Leroux, M. et al. Parameter estimation and control of an automatic balancing system. *Proc. SPIE* **13034** (2024). https://doi.org/10.1117/12.3013732

10. Yanyachi, Z. R. et al. Low-cost CubeSat attitude control test system. In *IEEE ARGENCON* (2022). https://doi.org/10.1109/ARGENCON55245.2022.9940099
11. Jovanovic, N., Pearce, J. M. & Praks, J. Low-cost open-source 3D printed air-bearing simulator. *J. Small Satellites* **8**, 859–880 (2019).
12. Rahnamai, K., Allison, J. & Peck, M. Low-cost CubeSat ADCS test platform. In *IEEE Aerospace Conf.* (2018). https://doi.org/10.1109/AERO.2018.8396419
13. Bryła, J. et al. 3D printed ADCS test platform for small satellites. *Int. J. Multiphysics* **15**, 397–408 (2021). https://doi.org/10.21152/1750-9548.15.4.397
14. Jamali, S. E., Okasha, M., Jallad, A., Dief, T. N. & Camps, A. Attitude Determination and Control System (ADCS) testbeds for CubeSats: A review. *IEEE Access* **14**, 68334–68356 (2026). https://doi.org/10.1109/ACCESS.2026.3685520
15. Piergentili, F. et al. CubeSat air-bearing V&V facility. ESA Workshop ADCSS (2024).
16. Ruf, O. et al. Automated CubeSat ADCS test environment. *CEAS Space J.* (2023). https://doi.org/10.1007/s12567-023-00523-x
17. Prado, J. et al. Three-axis air-bearing platform for ADCS simulation. *J. Appl. Res. Technol.* **3**, 222–237 (2005). https://doi.org/10.22201/icat.16656423.2005.3.03.563
18. Soh, G. S. et al. HIL ADCS test on air-bearing trolley. In *AIAA GNC Conf.* (2006). https://doi.org/10.2514/6.2006-6448
19. Menon, C. et al. Issues in free-flying robot testing. *Acta Astronaut.* **60**, 957–965 (2007). https://doi.org/10.1016/j.actaastro.2006.11.014
20. Schwartz, J. L. & Hall, C. D. System identification of spherical air-bearing simulator. *Adv. Astronaut. Sci.* **119**, 333–350 (2004).
21. Crowell, J. A. CubeSat ADCS testbed. M.S. thesis, MIT (2011).
22. Ustrzycki, A. et al. Spherical air-bearing simulator for nanosatellites. In *AIAA Small Satellites Conf.* (2011). https://doi.org/10.2514/6.2011-6272
23. Li, J. et al. CubeSat attitude control systems design. *J. Control Sci. Eng.* (2013). https://doi.org/10.1155/2013/657182
24. Kato, M. et al. FACE attitude control platform. *Trans. Jpn. Soc. Aeronaut. Space Sci.* **12**, Td_1–Td_6 (2014). https://doi.org/10.2322/tastj.12.Td_1
25. Post, M. A. et al. Magnetic field simulator for CubeSat testing. Zenodo (2014). https://doi.org/10.5281/zenodo.8347701
26. Kwan, T. H. et al. Air-bearing table for satellite simulation. In *IEEE ICIEA* (2015). https://doi.org/10.1109/ICIEA.2015.7334330
27. Tibbs, M. B. CubeSat ADCS testbed (6U). AFIT thesis (2015).
28. Gavrilovich, I. Robotic system for ADCS ground tests. PhD thesis, Univ. Montpellier (2016).
29. Bassett, C. CubeSat ADCS magnetic field test. AFIT thesis (2016).
30. Malekzadeh, M. et al. Sliding mode HIL attitude control. *Proc. IMechE G* **232**, 1944–1960 (2018). https://doi.org/10.1177/0954410017706992
31. Wolosik, M. J. CubeSat ADCS test development. M.Sc. thesis (2018).
32. Bahu, A. Magnetically actuated nanosatellite test facility. Univ. Bologna (2018).
33. da Silva, R. C. et al. Tabletop nanosatellite ADCS testbed. *J. Aerosp. Eng.* **32**, 04018122 (2019). https://doi.org/10.1061/(ASCE)AS.1943-5525.0000952
34. Cervettini, G. et al. Magnetic ADCS testbed. *IEEE Trans. Aerosp. Electron. Syst.* **57**, 1345–1350 (2021). https://doi.org/10.1109/TAES.2020.3040032
35. Papakonstantinou, C. et al. Low-cost air-bearing nano-CMG testbed. *Aerospace* **9**, 95 (2022). https://doi.org/10.3390/aerospace9020095
36. Pumacota, J. et al. MRAC-based CubeSat tracking testbed validation. In *IEEE INCAS* (2024). https://doi.org/10.1109/INCAS63820.2024.10798554
37. Chang, Y. et al. Air-bearing satellite control simulation. In *AIAA SciTech* (2025). https://doi.org/10.2514/6.2025-2827
38. Hoffmann, C. et al. Spherical air-bearing CubeSat testbed. In *AIAA SciTech* (2025). https://doi.org/10.2514/6.2025-1189
39. Ornati, F. et al. Vision-based attitude estimation for air-bearing simulators. *J. Spacecr. Rockets* (2024). https://doi.org/10.2514/1.A36073
40. Djebko, K. et al. AI-based satellite attitude controller (LeLaR). *IEEE Access* **14** (2026). https://doi.org/10.1109/ACCESS.2026.3678816