actuators

MDPI

Article

Observer-Based Hybrid Backstepping–Super-Twisting Control of a Twin Rotor MIMO System with Windowed Metaheuristic Gain Scheduling: Real-Time Tracking Experiments and Numerical Disturbance Analysis

Azeddine Beloufa [1], Abderrahmane Kacimi [2], Souaad Tahraoui [3], Abderrahmane Senoussaoui [4], Abdelbasset Azzouz [5], Mehdi Houari Zaid [6] and Jun-Jiat Tiang [7,*]

[1] Laboratory of Electrical Engineering and Renewable Energies (LGEER), Electronic Department, Hassiba Benbouali University, Chlef 02000, Algeria; a.beloufa@univ-chlef.dz
[2] Department of Instrumentation Maintenance, Institute of Maintenance and Industrial Safety, University of Oran 2, Oran 31000, Algeria; kacimi.abderrahmane@univ-oran2.dz
[3] Signals, Systems and Artificial Intelligence Laboratory (2SAIL), Electronic Department, Hassiba Benbouali University, Chlef 02000, Algeria; s.tahraoui@univ-chlef.dz
[4] Electrical Department, University of Mascara Mustapha Stambouli, Mascara 29000, Algeria; s.abderrahmane@univ-mascara.dz
[5] Laboratory of Electronics, Advanced Signal Processing, and Microwave (LESM), Department of Telecommunications, Saida University, Saida 20000, Algeria; abdelbasset.azzouz@univ-saida.dz
[6] Laboratoire AVCIS (Automatisation, Vision et Contrôle Intelligents des Systèmes), Département d'Automatique, Faculté de Génie Électrique, University of Science and Technology of Oran (USTO-MB), Oran 31000, Algeria; mehdihouari.zaid@univ-usto.dz
[7] Centre for Wireless Technology, CoE for Intelligent Network, Faculty of Artificial Intelligence and Engineering, Multimedia University, Cyberjaya 63100, Selangor, Malaysia
* Correspondence: jjtiang@mmu.edu.my

Abstract

Twin Rotor Multi-Input Multi-Output (TRMS) platforms combine strong aerodynamic cross-coupling, gravitational loading on the vertical axis, friction-dominated horizontal dynamics, and systematic mismatch between idealised models and laboratory hardware. The platform provides only two optical encoders, so the angular rates and the rotor states are unavailable for feedback. This paper presents an observer-based output-feedback architecture that addresses both difficulties. A high-gain observer built on the fully coupled six-state model, including the gyroscopic terms that the control design cannot retain, reconstructs the four unmeasured states from the two encoder angles. The reconstructed states drive a Hybrid Backstepping–Super-Twisting (B-STA) controller in which a second-order continuous sliding mode is embedded at the final recursive step through a composite surface. Because backstepping requires strict-feedback structure, which the centralised coupled model does not possess, the controller is synthesised on a decentralised design model and the residual coupling is rejected as a matched perturbation of the sliding variable. Closed-loop behaviour is analysed as a three-stage cascade covering observer error, sliding variable, and tracking error, yielding practical stability under bounded residuals with an explicit input-to-state gain. The residual bounds are evaluated numerically from the actuator saturation limit and the identified coefficients rather than assumed, and the resulting figures are shown to predict the marked difference in sliding-variable behaviour observed between the two axes. A second architecture applies a windowed Grey Wolf Optimiser (B-GWO) to the backstepping gains, in which each candidate is applied to the plant for a fixed test window, scored on its own accumulated integral of time-weighted absolute error, and followed by a settle window at the incumbent best. We prove that this windowing is a requirement rather than a convenience: a fitness evaluated at a single

Check for updates

Academic Editor: Guanghong Yang

Received: 18 July 2026
Revised: 17 August 2026
Accepted: 18 August 2026
Published: 20 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

sample is common to all candidates, cancels from the population ranking, and reduces the search to the minimiser of its own regularisation term. Both schemes are implemented on a physical TRMS through a Simulink Desktop Real-Time interface at a control period of 10 ms. On a 100 s experimental run, B-STA attains a pitch tracking error of 0.0318 rad RMS, 7.94% of the reference amplitude, and the lowest control energy on both axes among the strategies compared, reducing pitch control energy by 72.9% relative to a first-order Backstepping–Sliding Mode baseline recorded on the same interface. Numerical disturbance rejection tests on the fully coupled model confirm the mechanism: under a matched actuator step the super-twisting integrator state migrates to a new steady level that cancels the disturbance, driving the residual pitch error to 2×10^{-4} rad, whereas the same recursive law without the second-order injection retains a permanent offset of 0.28 rad.

Keywords: twin rotor MIMO system (TRMS); backstepping control; super-twisting algorithm; high-gain observer; output feedback; grey wolf optimiser (GWO); continuous sliding mode control; real-time experiment; second-order sliding mode

1. Introduction

1.1. Context

The Twin Rotor Multi-Input Multi-Output (TRMS) system is a laboratory benchmark for flight-oriented control of helicopters and Vertical Take-Off and Landing vehicles. Its difficulty is structural rather than incidental. The main and tail rotors generate thrust that scales quadratically with rotor speed, the vertical axis carries a gravitational restoring torque that the horizontal axis does not, the two axes exchange reaction and gyroscopic torques, and the mechanical assembly contributes friction and cable tension that no manufacturer model captures [1]. Linear designs such as the Linear Quadratic Regulator [2] and Model Predictive Control [3] perform adequately near a design equilibrium and degrade once the trajectory leaves the linearisation neighbourhood.

Recursive nonlinear designs, principally Backstepping and Sliding Mode Control, are the standard response [4,5]. Each carries a known cost. Backstepping cancels modelled dynamics exactly and is therefore only as good as the model [5]. First-order sliding mode replaces that dependence with a discontinuous injection, and the injection appears directly in the actuator command. On a TRMS this is not a cosmetic issue: the voltage switches at the sampling rate, the motors dissipate the resulting current ripple as heat, and the shafts and bearings absorb the vibratory load.

1.2. Current State of the Field

Three lines of work bear directly on the present study, and it is useful to state plainly what each has already established.

Second-order sliding mode on the TRMS itself. The Super-Twisting Algorithm of Levant [6], given strict Lyapunov functions by Moreno and Osorio [7], moves the discontinuity into the derivative of an auxiliary state, so that the applied control is continuous while finite-time convergence of the sliding variable is preserved. Mondal and Mahanta applied an adaptive second-order sliding mode controller to a TRMS and demonstrated the expected reduction in chattering [8]. Super-twisting designs have also been applied to robotic manipulators [9] and to doubly fed induction generators in wind energy conversion [10]. Second-order sliding mode on the TRMS is therefore not new, and this paper does not claim it as such.

Backstepping combined with super-twisting. The combination has been reported for deep-sea hydraulic manipulators with an observer in the loop [11], and related backstepping and sliding mode hybrids appear in hydraulic drives [12], permanent magnet synchronous machines [13], coupled unmanned aerial vehicles [14], aircraft lateral control [15], wind energy conversion [16], microgrids [17], multiphase drives [18], and linear synchronous motors [19]. The structural idea is established. What remains open on the TRMS is quantitative: how much actuator load and tracking error a continuous second-order injection actually removes on physical hardware, measured against a first-order baseline recorded under an identical interface.

Online metaheuristic gain tuning. Metaheuristic PID design [20], multi-objective genetic tuning [21], fractional-order schemes [22,23], particle-swarm-assisted control [24], and optimised backstepping for electro-hydraulic actuators [25] have all been reported, and the Grey Wolf Optimiser [26] has been used for controller tuning [27] and multi-objective search [28]. In our own prior work an online Rate-Constrained Sequential Quantum Black Hole Optimiser was applied to this platform [29], and a comparison of adaptive tuning functions against metaheuristic gain optimisation on the same rig is reported in [30]. The latter is the closest antecedent to the scheduler developed in Section 6 and differs from it in the respect that matters here: the candidate fitness there is formed at the sampling instant rather than accumulated over an evaluation window, which is precisely the configuration Proposition 1 shows to be degenerate. Almost all of this literature evaluates candidates offline against a model. The question of what a population-based search can and cannot do when the fitness must be measured on a running plant has received far less attention, and it is where the present paper makes a specific contribution.

1.3. The Gap Addressed Here

Two issues are unresolved by the work above, and they interact.

The first is observability. A TRMS provides two optical encoders and nothing else. The angular rates and both rotor states are unmeasured. Backstepping on a third-order chain requires all three states of that chain, and the super-twisting surface requires them as well. Studies reporting backstepping or sliding mode results on this platform frequently do not state how the missing states were obtained, yet the answer determines what the stability analysis actually proves. A controller analysed as full-state feedback but implemented on estimates is not the system that was analysed.

The second is structural. The coupled six-state TRMS model is not in strict-feedback form. The pitch chain contains a yaw state at the wrong level, so recursive backstepping cannot be applied to the centralised model at all. Every backstepping result on this platform therefore rests on a decentralised decomposition, but the decomposition is usually presented as a simplification rather than as the structural requirement it is, and the residual coupling is rarely characterised.

This paper addresses both. The observer uses the fully coupled model, including the gyroscopic terms, so the coupling is compensated where compensation is structurally permitted. The controller is decentralised because strict-feedback form demands it, and the residual coupling is shown to be matched with respect to the sliding variable, which is what makes it rejectable by the super-twisting injection rather than merely bounded.

1.4. Contributions

1. An observer-based output-feedback architecture for the TRMS in which a high-gain observer on the fully coupled six-state model reconstructs the four unmeasured states from two encoder angles and drives a Hybrid Backstepping–Super-Twisting controller built on a composite surface combining all three backstepping errors.

2. A three-stage cascade stability analysis covering observer error, sliding variable, and tracking error. The reduced tracking dynamics are shown to be input-to-state stable with respect to the sliding variable with an explicit gain, so that practical stability follows without assuming ideal sliding, and the residual bound is derived from the coupling coefficients and the actuator saturation limit rather than assumed.
3. A numerical evaluation of that residual bound on both axes, which shows the perturbation entering the yaw sliding variable to be approximately six times that entering the pitch sliding variable, and which accounts for the difference in sliding-variable behaviour observed between the two axes in both simulation and experiment.
4. An explicit justification of the decentralised decomposition on strict-feedback grounds, together with a demonstration that the neglected coupling is matched with respect to the composite sliding variable and therefore lies in the channel the super-twisting term can reject.
5. A windowed Grey Wolf gain scheduler in which candidates are applied to the plant and scored on their own accumulated error, accompanied by a proof that this windowing is necessary: any fitness of the form $g(e(k)) + h(\boldsymbol{\theta}_i)$ evaluated before the candidate acts on the plant reduces to $\arg\min h$, independent of closed-loop behaviour.
6. Real-time experimental validation on a physical TRMS reporting normalised tracking error, control energy, and total variation in the applied voltage against a first-order baseline recorded on the same interface, together with numerical disturbance rejection tests on the fully coupled model under matched load, parametric, and aerodynamic perturbations.

1.5. Organisation

Section 2 presents the fully coupled TRMS model and the decentralised design model, and states precisely what the decomposition discards. Section 3 presents the observer. Section 4 develops the B-STA control law. Section 5 gives the cascade stability analysis. Section 6 develops the windowed B-GWO scheduler and the degeneracy result. Section 7 reports nominal and disturbance rejection simulations. Section 8 reports the real-time experiments. Section 9 concludes.

2. TRMS Model and Design Decomposition

2.1. Plant Description and Coordinate Convention

The experimental platform is a two-degree-of-freedom Twin Rotor MIMO System. A beam is mounted on a vertical tower and is free to rotate in the vertical plane through the pitch angle ψ and in the horizontal plane through the yaw angle ϕ. The main rotor at one beam extremity generates thrust acting principally on pitch; the tail rotor at the other extremity acts principally on yaw. A counterweight arm provides static balance. Figure 1 defines the coordinate system, the positive senses of ψ and ϕ, the thrust vectors of the two rotors, the gravitational moment, and the direction of the gyroscopic coupling torque.

2.2. Fully Coupled Nonlinear Model

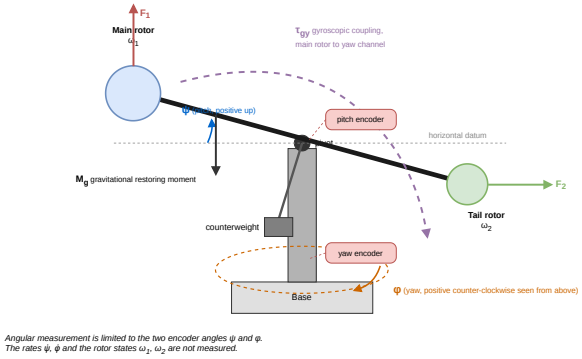
Let the state vector be

$$\mathbf{x} = [x_1, x_2, x_3, x_4, x_5, x_6]^\top = [\psi, \dot{\psi}, \omega_1, \phi, \dot{\phi}, \omega_2]^\top, \tag{1}$$

where ω_1 and ω_2 are the main and tail rotor states, and let $\mathbf{u} = [U_1, U_2]^\top$ be the DC voltages applied to the two motors.

Both notations are used in what follows, and the distinction is deliberate rather than incidental. The indexed form is retained wherever the argument concerns the position of a state within its chain, because two of the results below are statements of exactly that

kind. Section 2.3 establishes that the coupled model is not in strict-feedback form because the equation for $\dot{x}_2$ contains x_5, a state belonging to the other chain and entering at the wrong level; written in indexed form this is visibly a violation of triangular ordering, whereas the same fact expressed as $\dot{\phi}$ entering the $\ddot{\psi}$ equation reads only as a physical remark about gyroscopic coupling. The observability indices $\mu_1 = \mu_2 = 3$ of Section 3 and the block scaling built on them refer likewise to the depth of each chain. The physical symbols are used where the discussion concerns the platform itself, in the figures and in the experimental sections, and (1) is the correspondence between the two.



(**a**)



(**b**)

Figure 1. The Twin Rotor MIMO System. (**a**) Coordinate schematic. The pitch angle ψ is defined about the horizontal axis and the yaw angle ϕ about the vertical axis, with positive senses indicated. F_1 and F_2 denote the main and tail rotor thrust vectors, M_g the gravitational restoring moment acting on the pitch axis, and τ_{gy} the gyroscopic coupling torque transmitted from the main rotor to the yaw channel. Angular measurement is provided by two incremental optical encoders mounted on the pitch and yaw axes; no rate or rotor state sensing is available. (**b**) The laboratory platform on which the experiments of Section 8 were performed, photographed by the authors.

The complete nonlinear model of the platform is

$$\begin{cases} \dot{x}_1 = x_2 \\ \dot{x}_2 = \dfrac{1}{I_1}\Big(a_1 x_3^2 + b_1 x_3 - M_g \sin x_1 - B_{1\psi} x_2 - B_{2\psi} x_1 - K_{gy} x_5 \cos x_1 \big(a_1 x_3^2 + b_1 x_3\big)\Big) \\ \dot{x}_3 = \dfrac{-T_{10}\, x_3 + K_1 U_1}{T_{11}} \\ \dot{x}_4 = x_5 \\ \dot{x}_5 = \dfrac{1}{I_2}\Big(a_2 x_6^2 + b_2 x_6 - B_{1\phi} x_5 - B_{2\phi} x_4 - 1.75 K_c \big(a_1 x_3^2 + b_1 x_3\big)\Big) \\ \dot{x}_6 = \dfrac{-T_{20}\, x_6 + K_2 U_2}{T_{21}} \end{cases} \tag{2}$$

The two coupling terms have different physical origins and different structure. The term $K_{gy} x_5 \cos x_1 (a_1 x_3^2 + b_1 x_3)$ in the pitch equation is gyroscopic: it is proportional to the yaw rate and to the main rotor thrust, and it vanishes when either is zero. The term $1.75 K_c (a_1 x_3^2 + b_1 x_3)$ in the yaw equation is the reaction torque transmitted to the horizontal axis by the main rotor, and it depends on the main rotor state alone. Numerical values follow the manufacturer specification [31] together with the identification reported in [32], and are listed in Table 1. The same values are used in the observer, in both controllers, and in the simulation model, so that no discrepancy exists between the model on which the estimator is built and the model on which the control laws are synthesised.

Table 1. Physical parameters of the TRMS platform.

Parameter	Value	Parameter	Value
I_1	6.8×10^{-2} kg·m^2	I_2	2.0×10^{-2} kg·m^2
a_1	0.0135	a_2	0.02
b_1	0.0924	b_2	0.09
M_g	0.32 N·m	$B_{1\psi}$	6.0×10^{-3} N·m·s/rad
$B_{2\psi}$	1.0×10^{-3} N·m/rad	$B_{1\phi}$	1.0×10^{-1} N·m·s/rad
$B_{2\phi}$	1.0×10^{-2} N·m/rad	K_{gy}	0.05 s/rad
K_c	-0.2	K_1	1.1
K_2	0.8	T_{11}	1.1
T_{10}	1.0	T_{21}	1.0
T_{20}	1.0		

2.3. Why the Design Model Must Be Decentralised

Recursive backstepping is applicable only to systems in strict-feedback form, in which the state entering each equation of the chain is the state advanced by the next equation, and no state from a later stage appears earlier. Inspection of (2) shows that the pitch chain violates this condition. The equation for $\dot{x}_2$ contains x_5, which belongs to the yaw chain and is not the virtual control of the pitch chain at that level. A virtual control designed at the second pitch step would therefore have to cancel a term over which it has no authority, and the recursion cannot close. The same argument applies to the yaw chain, whose $\dot{x}_5$ equation contains x_3.

The centralised coupled model is consequently not a candidate for backstepping synthesis at all. This is a structural obstruction and not a matter of modelling convenience. Two responses are available. The first is to seek an exact input–output decoupling or a feedback linearising transformation that restores triangularity. This requires an invertible decoupling matrix over the whole operating envelope and inherits full sensitivity to the coupling coefficients, which on this platform are among the least reliably identified parameters. The second, adopted here, is to synthesise the controller on a decentralised design model that is strict-feedback by construction, and to treat what the decomposition discards as a perturbation to be rejected rather than cancelled. The advantage of the second route is that rejection does not require the coupling coefficients to be known accurately, only bounded.

We emphasise that the coupling is discarded only from the *control design model*. It is retained in full in the observer of Section 3, and it is present in the physical plant at all times.

2.4. Decentralised Design Model

Removing the coupling terms from (2) yields two independent third-order strict-feedback chains. The pitch design model is

$$\begin{cases} \dot{x}_1 = x_2 \\ \dot{x}_2 = \dfrac{1}{I_1}\left(a_1 x_3^2 + b_1 x_3 - M_g \sin x_1 - B_{1\psi} x_2 - B_{2\psi} x_1\right) \\ \dot{x}_3 = \dfrac{-T_{10}\, x_3 + K_1\, U_1}{T_{11}} \end{cases} \tag{3}$$

and the yaw design model is

$$\begin{cases} \dot{x}_4 = x_5 \\ \dot{x}_5 = \dfrac{1}{I_2}\left(a_2 x_6^2 + b_2 x_6 - B_{1\phi} x_5 - B_{2\phi} x_4\right) \\ \dot{x}_6 = \dfrac{-T_{20}\, x_6 + K_2\, U_2}{T_{21}} \end{cases} \tag{4}$$

The discarded terms are collected as

$$d_\psi(\mathbf{x}) \triangleq -\frac{K_{gy}}{I_1}\, x_5 \cos x_1 \left(a_1 x_3^2 + b_1 x_3\right), \qquad d_\phi(\mathbf{x}) \triangleq -\frac{1.75 K_c}{I_2}\left(a_1 x_3^2 + b_1 x_3\right). \tag{5}$$

A structural asymmetry between the two channels governs much of what follows. The pitch dynamics carry the gravitational restoring torque $M_g \sin x_1$, which acts as a persistent state-dependent load; the yaw dynamics carry no such term but a viscous friction coefficient $B_{1\phi}$ roughly seventeen times larger than $B_{1\psi}$, acting on a moment of inertia roughly one third as large. A second asymmetry is visible in (5) itself and is quantified in Section 2.5: the gyroscopic term acting on pitch is proportional to the yaw rate and therefore vanishes whenever the horizontal axis is stationary, whereas the reaction torque acting on yaw depends on the main rotor state alone and is present whenever the main rotor turns.

2.5. Bounds on the Neglected Coupling

Assumption 1 in Section 5 requires the residual perturbation and its derivative to be bounded. Because the platform is physically constrained, these bounds can be derived rather than postulated.

The pitch and yaw angles are limited by mechanical stops to $|x_1| \le \psi_{\max}$ and $|x_4| \le \phi_{\max}$. The applied voltages are saturated in software at $|U_i| \le U_{\max} = 2.5\,\mathrm{V}$. The rotor states obey stable first-order dynamics driven by these voltages, so from (2),

$$|x_3| \le \frac{K_1 U_{\max}}{T_{10}} \triangleq \omega_{1,\max}, \qquad |x_6| \le \frac{K_2 U_{\max}}{T_{20}} \triangleq \omega_{2,\max} \tag{6}$$

in steady state, and the angular rates are bounded on the compact operating set by boundedness of the right-hand sides of (2). Writing $\Theta_{\max} \triangleq a_1 \omega_{1,\max}^2 + b_1 \omega_{1,\max}$ for the maximum main rotor thrust term, the coupling magnitudes obey

$$|d_\psi| \le \frac{K_{gy}}{I_1}\, \dot{\phi}_{\max}\, \Theta_{\max}, \qquad |d_\phi| \le \frac{1.75 |K_c|}{I_2}\, \Theta_{\max}. \tag{7}$$

Differentiating (5) and bounding each term likewise gives finite bounds on $|\dot{d}_\psi|$ and $|\dot{d}_\phi|$, in which $\dot{x}_3$ and $\dot{x}_6$ are bounded through (2) and the saturation limit. The derivative bound therefore depends explicitly on $U_{\max}$. This is a consequence of the design rather than a deficiency of the analysis: it is precisely because the actuator is saturated that the perturbation derivative admits a bound at all.

Numerical Evaluation

The bounds above are evaluated with the coefficients of Table 1, the saturation limit $U_{\max} = 2.5\,\mathrm{V}$, and the angular rates commanded by the reference trajectories of Section 8, which have period 12.57 s and amplitudes 0.40 rad on pitch and 1.00 rad on yaw, giving $\dot{\psi}_{\max} = 0.20\,\mathrm{rad/s}$ and $\dot{\phi}_{\max} = 0.50\,\mathrm{rad/s}$. From (6), $\omega_{1,\max} = 2.75$ and $\omega_{2,\max} = 2.00$, hence $\Theta_{\max} = 0.356$.

The perturbation does not enter the sliding variable at unit gain. Substituting (5) into the derivative of (24) through the route identified in Section 4.3 shows that d_ψ reaches $\dot{s}$ through the $c_2\dot{\hat{x}}_2$ term of φ_p and through the $-(k_1+k_2)\dot{\hat{x}}_2$ contribution to $\dot{\alpha}_2$, so that

$$\delta_\psi \simeq (c_2 + k_1 + k_2)\, d_\psi, \qquad \delta_\phi \simeq (c_5 + k_4 + k_5)\, d_\phi, \tag{8}$$

with transfer gains of 17 on pitch and 43 on yaw for the controller gains adopted in Section 4. The resulting figures are collected in Table 2.

Table 2. Numerical evaluation of the residual perturbation bounds of Assumption 1, computed from Table 1, the saturation limit $U_{\max} = 2.5\,\mathrm{V}$ and the commanded reference rates. All quantities are envelope bounds and are attained only if every contributing term reaches its individual maximum simultaneously.

Quantity	Pitch	Yaw	Source
Rotor state bound $\omega_{i,\max}$	2.75	2.00	(6)
Thrust term $\Theta_{\max}$	0.356		definition
Coupling $\lvert d\rvert$ (rad/s^2)	0.131	6.23	(7)
Transfer gain into $\dot{s}$	17	43	(8)
Δ_0 (rad/s^3)	2.23	268	(39)
Coupling rate $\lvert\dot{d}\rvert$ (rad/s^3)	6.16	14.6	derivative of (5)
Δ_1 (rad/s^4)	105	627	(39)

Two features of Table 2 govern the interpretation of the stability result and are stated here so that Section 5 can refer to them directly.

First, these are envelope bounds. They are composed of the maximum rotor thrust, the maximum yaw rate and the maximum angular acceleration, a combination that no trajectory generated by the references of Section 8 attains. They are the correct quantities for a worst-case guarantee and a substantial overestimate of the perturbation actually experienced on the recorded runs.

Second, and independently of that conservatism, the ratio between the two axes is a property of the platform rather than of the bounding procedure. The yaw figure exceeds the pitch figure by a factor of approximately six in Δ_1 and by two orders of magnitude in Δ_0. Two mechanisms compound. The reaction-torque coefficient $1.75|K_c|/I_2 = 17.5$ is an order of magnitude larger than the gyroscopic coefficient $K_{gy}\dot{\phi}_{\max}/I_1 = 0.37$, and the yaw transfer gain of 43 is two and a half times the pitch value because the yaw backstepping gains are set higher for the reasons given in Section 4. The choice that improves yaw tracking through faster error convergence therefore simultaneously amplifies the residual coupling entering the yaw sliding variable. Section 5.2 draws the consequence, and Section 7.3 shows that the predicted ordering is the one observed.

3. High-Gain Observer for the Unmeasured States

3.1. Measurement Limitation

The TRMS is instrumented with two incremental optical encoders, one on each axis. The measured output is therefore

$$\mathbf{y} = \mathbf{C}_0\mathbf{x} = [x_1,\ x_4]^\top, \qquad \mathbf{C}_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \tag{9}$$

The angular rates x_2 and x_5 and both rotor states x_3 and x_6 are unmeasured. Since the control laws of Section 4 require all six states, and in particular require x_3 and x_6 inside the input gain, an observer is a structural necessity rather than an implementation convenience.

Numerical differentiation of the encoder signals is not a viable substitute: the encoder quantises the measured angle, and the control law depends on quantities as high as the third derivative of the reference and the first derivative of the rotor state.

3.2. Observer Structure

We employ the high-gain observer of Bornard and Hammouri [33], in the form previously developed and experimentally validated by the present authors on this platform [32]. For a MIMO nonlinear system $\dot{\mathbf{x}} = f(\mathbf{x}) + \mathbf{B}\mathbf{u}$, $\mathbf{y} = h(\mathbf{x})$, the observer dynamics are

$$\dot{\hat{\mathbf{x}}} = f(\hat{\mathbf{x}}) + \mathbf{B}\mathbf{u} + \left(\frac{\partial \varphi(\hat{\mathbf{x}})}{\partial \hat{\mathbf{x}}}\right)^{-1} \mathbf{\Lambda}^{-1}(\theta)\, \mathbf{L}\left(\mathbf{y} - \mathbf{C}_0 \hat{\mathbf{x}}\right), \tag{10}$$

where φ is the observability map, $\partial\varphi/\partial\hat{\mathbf{x}}$ its Jacobian, $\mathbf{L}$ a constant correction gain, and $\mathbf{\Lambda}^{-1}(\theta)$ a block-diagonal scaling matrix parameterised by the tuning gain θ. The correction term scaled by $\mathbf{\Lambda}^{-1}$ dominates the system nonlinearities, forcing the estimation error dynamics to behave as a stable linear system and yielding exponential convergence under standard observability conditions.

Two features of the present implementation are essential to the argument of this paper and are stated explicitly.

First, the vector field $f(\hat{\mathbf{x}})$ in (10) is that of the *fully coupled* model (2), including both the K_{gy} and K_c terms. The observability map and its Jacobian are constructed accordingly, so that the off-diagonal entries $\partial L_f^2 h_1/\partial \hat{x}_4$ and $\partial L_f^2 h_2/\partial \hat{x}_5$ carry the coupling explicitly. The estimator is not subject to the strict-feedback restriction that forces the controller to be decentralised, and it therefore compensates the cross-coupling where compensation is structurally permitted. The complete derivation of the observability map, the Jacobian entries, and the associated convergence proof are given in [32] and are not reproduced here. A related observer-based backstepping architecture for the same platform, using an interval type-2 fuzzy estimator in place of the high-gain structure, is reported in [34].

Second, the observer is applied with $n = 6$, $p = 2$, and observability indices $\mu_1 = \mu_2 = 3$, giving the block scaling

$$\mathbf{\Lambda}^{-1}(\theta) = \mathrm{blkdiag}\big(\mathbf{\Delta}_1(\theta_1),\, \mathbf{\Delta}_2(\theta_2)\big), \qquad \mathbf{\Delta}_i(\theta_i) = \mathrm{diag}\big(1,\, \theta_i^{-1},\, \theta_i^{-2}\big), \tag{11}$$

with correction gains $\mathbf{L}_i = [3\theta_i,\, 3\theta_i^2,\, \theta_i^3]^\top$ placing the nominal error poles.

3.3. Gain Selection and Its Consequence for the Analysis

The gain values used in the experiments reported here are $\theta_1 = \theta_2 = \theta^\star = 1$, so that $\mathbf{\Lambda}^{-1}$ reduces to the identity and the nominal error poles of each block are placed at -1 by $\mathbf{L}_i = [3,\, 3,\, 1]^\top$.

At this gain the observer does not operate in the asymptotic high-gain regime, so the classical separation argument based on timescale decoupling is not invoked, and neither is the peaking-containment analysis that accompanies it. The observer–controller interconnection is instead analysed directly as the first stage of the cascade in Section 5. Two properties make this tractable. The applied voltages are globally bounded by the software saturation at ± 2.5 V, so a transient in the estimated states cannot propagate without limit into the plant. And the estimation error enters the control law only through terms that are Lipschitz on the compact operating set, so it appears in the sliding dynamics as an additive matched perturbation of bounded magnitude, which is exactly the class of perturbation the super-twisting injection is designed to reject.

The choice $\theta^\star = 1$ is deliberate and reflects the quantisation of the encoder signal. Raising θ scales the correction gains as θ^3, which accelerates convergence of the estimate at

the cost of amplifying the quantisation step of the encoder into the reconstructed rate and rotor states. On this platform the reconstructed states drive an input gain that is inverted in (29), so estimator noise enters the actuator command directly, and the moderate gain was preferred.

Writing $\tilde{\mathbf{x}} \triangleq \hat{\mathbf{x}} - \mathbf{x}$, exponential convergence gives constants $c_o, \sigma_o > 0$ such that

$$\|\tilde{\mathbf{x}}(t)\| \leq c_o \|\tilde{\mathbf{x}}(0)\| e^{-\sigma_o t} + \kappa_o, \tag{12}$$

where κ_o collects the effect of measurement quantisation and unmodelled dynamics on the estimate. The residual bound κ_o contributes additively to the perturbation bounds Δ_0 and Δ_1 of Assumption 1, over and above the coupling contribution evaluated in Section 2.5.

3.4. Implemented Control Uses Estimated States

All control expressions in Section 4 are written in terms of the estimated states $\hat{x}_i$, since these are the quantities available to the implemented law. Where the true state appears in an analytical statement it is written x_i without a hat. Figure 2 shows the observer explicitly in the loop.

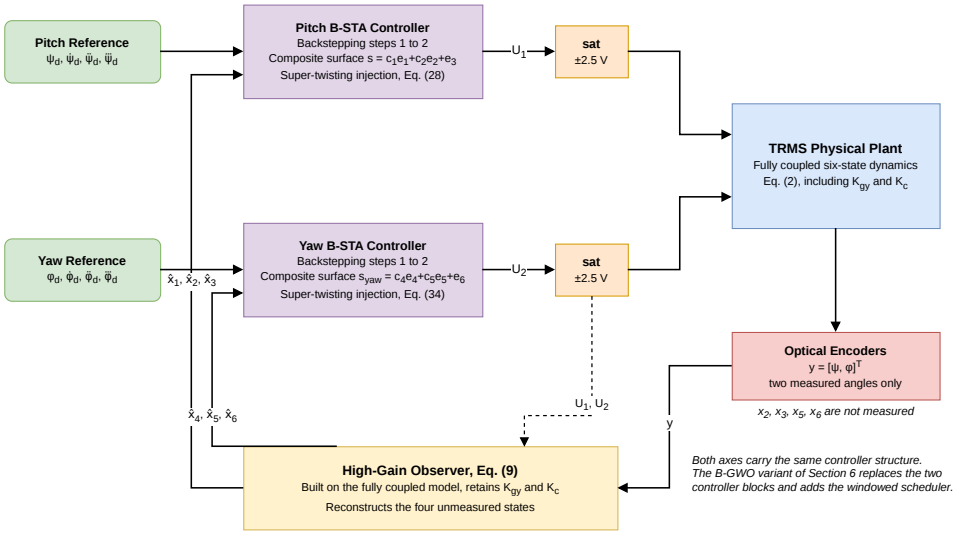


Figure 2. Observer-based decentralised control architecture. The two measured encoder angles $\mathbf{y} = [\psi, \phi]^\top$ drive a high-gain observer constructed on the fully coupled six-state model. The estimated pitch states $[\hat{x}_1, \hat{x}_2, \hat{x}_3]$ and yaw states $[\hat{x}_4, \hat{x}_5, \hat{x}_6]$ are routed to their respective controllers. Both axes carry the same controller structure. Applied voltages are saturated at ± 2.5 V before reaching the plant.

4. Hybrid Backstepping–Super-Twisting Controller

The architecture is shown in Figure 2. The two encoder angles drive the coupled-model observer, which supplies all six estimated states. The pitch and yaw controllers are synthesised independently on the strict-feedback design models (3) and (4) and operate simultaneously on the physical plant, where the coupling remains present.

4.1. Step 1: Position Tracking Error

Define the pitch tracking error

$$e_1 \triangleq \hat{x}_1 - x_{1d}, \tag{13}$$

where $x_{1d}(t)$ is a sufficiently smooth reference with bounded derivatives to third order. The first virtual control is

$$\alpha_1 \triangleq -k_1 e_1 + \dot{x}_{1d}, \qquad k_1 > 0, \tag{14}$$

and with the velocity error

$$e_2 \triangleq \hat{x}_2 - \alpha_1 \tag{15}$$

the first error equation becomes

$$\dot{e}_1 = -k_1 e_1 + e_2. \tag{16}$$

Differentiating (14),

$$\dot{\alpha}_1 = -k_1(\hat{x}_2 - \dot{x}_{1d}) + \ddot{x}_{1d}. \tag{17}$$

4.2. Step 2: Velocity Tracking Error

The pitch acceleration predicted by the design model is

$$\dot{\hat{x}}_2 = \frac{1}{I_1}\Big(a_1 \hat{x}_3^2 + b_1 \hat{x}_3 - M_g \sin \hat{x}_1 - B_{1\psi}\hat{x}_2 - B_{2\psi}\hat{x}_1\Big). \tag{18}$$

It is convenient to split this into the part controllable through the rotor state and the gravitational stiffness part,

$$\mathcal{E} \triangleq \frac{1}{I_1}\Big(a_1 \hat{x}_3^2 + b_1 \hat{x}_3 - B_{1\psi}\hat{x}_2\Big), \qquad \mathcal{D}(\hat{x}_1) \triangleq \frac{1}{I_1}\big(M_g \sin \hat{x}_1 + B_{2\psi}\hat{x}_1\big), \tag{19}$$

so that $\dot{\hat{x}}_2 = \mathcal{E} - \mathcal{D}(\hat{x}_1)$. The second virtual control is chosen as

$$\alpha_2 \triangleq -k_2 e_2 - e_1 + \mathcal{D}(\hat{x}_1) + \dot{\alpha}_1, \qquad k_2 > 0, \tag{20}$$

which on substituting (17) becomes

$$\alpha_2 = -k_2 e_2 - e_1 + \mathcal{D}(\hat{x}_1) - k_1 \hat{x}_2 + k_1 \dot{x}_{1d} + \ddot{x}_{1d}. \tag{21}$$

Defining the third error

$$e_3 \triangleq \mathcal{E} - \alpha_2 \tag{22}$$

gives the second error equation

$$\dot{e}_2 = -k_2 e_2 - e_1 + e_3. \tag{23}$$

4.3. Step 3: Composite Surface and Super-Twisting Injection

Rather than driving e_3 to zero by a third backstepping step, which would again require exact cancellation, we introduce the composite sliding surface

$$s \triangleq c_1 e_1 + c_2 e_2 + e_3, \qquad c_1, c_2 > 0. \tag{24}$$

The surface coefficients carry the units required to render s dimensionally consistent, since e_1, e_2 and e_3 have units of rad, rad/s and rad/s^2 respectively.

Combining (24) with (16) and (23) and eliminating e_3 gives the exact relation

$$\begin{bmatrix} \dot{e}_1 \\ \dot{e}_2 \end{bmatrix} = \underbrace{\begin{bmatrix} -k_1 & 1 \\ -(1+c_1) & -(k_2+c_2) \end{bmatrix}}_{\mathbf{A}_{\mathrm{cl}}} \begin{bmatrix} e_1 \\ e_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} s, \tag{25}$$

which holds for all t and not only on the sliding set. This relation is the basis of the cascade analysis in Section 5.

Choice of the composite surface.

The composite construction is not merely a design preference. From (25) and (23), the neglected coupling d_ψ of (5) enters $\dot{s}$ through the $c_2\dot{e}_2$ term and through the $-(k_1+k_2)\dot{\hat{x}}_2$ contribution to $\dot{\alpha}_2$, which is the same place the control input enters through $g(\hat{\mathbf{x}})U_1$. The coupling is therefore *unmatched* with respect to the original state chain, where it acts at $\dot{x}_2$ while the input acts at $\dot{x}_3$, but *matched* with respect to the composite sliding variable. This is what places it in the channel the super-twisting term can reject, and it is the reason a surface built on e_3 alone would not suffice. The same two routes fix the transfer gain $(c_2+k_1+k_2)$ used in (8).

Derivative of the auxiliary variable.

Differentiating $\mathcal{E}$ and substituting the motor dynamics isolates the control channel,

$$\dot{\mathcal{E}} = \underbrace{\frac{(2a_1\hat{x}_3+b_1)(-T_{10}\hat{x}_3)}{I_1T_{11}} - \frac{B_{1\psi}\dot{\hat{x}}_2}{I_1}}_{\dot{\mathcal{E}}_{\mathrm{nom}}(\hat{\mathbf{x}})} + \underbrace{\frac{K_1(2a_1\hat{x}_3+b_1)}{I_1T_{11}}}_{g(\hat{\mathbf{x}})}U_1. \tag{26}$$

The input gain $g(\hat{\mathbf{x}})$ vanishes only at $\hat{x}_3 = -b_1/(2a_1) \approx -3.42$, which lies outside the physical main rotor envelope established in (6). Because $\hat{x}_3$ is an estimate rather than a measurement, a numerical safeguard $|g| \geq \epsilon_g = 10^{-4}$ is enforced in software, protecting against transient excursions of the estimate during observer convergence and not only against physically attainable operating points.

Derivative of the second virtual control.

Expressing (21) in explicit state form and differentiating term by term yields

$$\begin{aligned}\dot{\alpha}_2 = {} & -(k_1+k_2)\,\dot{\hat{x}}_2 + \hat{x}_2\left(-k_1k_2 - 1 + \frac{M_g\cos\hat{x}_1 + B_{2\psi}}{I_1}\right) \\ & + (k_1k_2+1)\,\dot{x}_{1d} + (k_1+k_2)\,\ddot{x}_{1d} + \dddot{x}_{1d}.\end{aligned} \tag{27}$$

Input-affine sliding dynamics and control law.

Substituting (26) into the derivative of (24) gives

$$\dot{s} = \varphi_p(\hat{\mathbf{x}},t) + g(\hat{\mathbf{x}})\,U_1, \qquad \varphi_p = c_1(\hat{x}_2 - \dot{x}_{1d}) + c_2\big(\dot{\hat{x}}_2 + k_1(\hat{x}_2 - \dot{x}_{1d}) - \ddot{x}_{1d}\big) + \dot{\mathcal{E}}_{\mathrm{nom}} - \dot{\alpha}_2. \tag{28}$$

The applied pitch voltage combines an equivalent part cancelling the nominal drift with a super-twisting part,

$$U_1 = \mathrm{sat}_{2.5}\left\{\frac{1}{g(\hat{\mathbf{x}})}\Big[-\varphi_p(\hat{\mathbf{x}},t) - \lambda_1|s|^{1/2}\,\mathrm{sgn}(s) + w\Big]\right\}, \tag{29}$$

$$\dot{w} = -\lambda_2\,\mathrm{sgn}(s), \qquad w(0) = 0, \qquad \lambda_1,\lambda_2 > 0, \tag{30}$$

where $\mathrm{sat}_{2.5}\{\cdot\}$ denotes symmetric saturation at ± 2.5 V. Under exact cancellation the sliding dynamics reduce to the ideal super-twisting form

$$\dot{s} = -\lambda_1|s|^{1/2}\,\mathrm{sgn}(s) + w, \qquad \dot{w} = -\lambda_2\,\mathrm{sgn}(s). \tag{31}$$

The discontinuity acts on $\dot{w}$ rather than on U_1. Since w is absolutely continuous and $|s|^{1/2}\,\mathrm{sgn}(s)$ is continuous and vanishes at $s = 0$, the unsaturated command is a continuous function of time.

In discrete time at the implementation sample period $T_s = 0.01\,\mathrm{s}$, the auxiliary state is updated by forward Euler integration,

$$w(k+1) = w(k) - T_s \lambda_2 \operatorname{sgn}\big(s(k)\big), \tag{32}$$

and clamped at $|w| \le w_{\max} = 50$ to limit windup while the command is saturated.

4.4. Yaw Subsystem

The yaw synthesis mirrors the pitch derivation on the design model (4). With errors $e_4 = \hat{x}_4 - x_{4d}$ and $e_5 = \hat{x}_5 - \alpha_4$, virtual controls

$$\alpha_4 = -k_4 e_4 + \dot{x}_{4d}, \qquad \alpha_5 = -k_5 e_5 - e_4 + \frac{B_{2\phi}\hat{x}_4}{I_2} - k_4 \hat{x}_5 + k_4 \dot{x}_{4d} + \ddot{x}_{4d}, \tag{33}$$

auxiliary variable $\mathcal{F} = (a_2 \hat{x}_6^2 + b_2 \hat{x}_6 - B_{1\phi}\hat{x}_5)/I_2$, third error $e_6 = \mathcal{F} - \alpha_5$, and composite surface

$$s_{\mathrm{yaw}} \triangleq c_4 e_4 + c_5 e_5 + e_6, \tag{34}$$

the control law is

$$U_2 = \mathrm{sat}_{2.5}\left\{ \frac{1}{g_{\mathrm{yaw}}(\hat{\mathbf{x}})} \left[-\varphi_y(\hat{\mathbf{x}}, t) - \lambda_3 |s_{\mathrm{yaw}}|^{1/2} \operatorname{sgn}(s_{\mathrm{yaw}}) + w_{\mathrm{yaw}} \right] \right\}, \tag{35}$$

$$\dot{w}_{\mathrm{yaw}} = -\lambda_4 \operatorname{sgn}(s_{\mathrm{yaw}}), \qquad g_{\mathrm{yaw}}(\hat{\mathbf{x}}) = \frac{K_2(2a_2\hat{x}_6 + b_2)}{I_2 T_{21}}. \tag{36}$$

The virtual control derivative $\dot{\alpha}_5$ follows (27) with the gravitational Jacobian term absent, consistent with the horizontal axis carrying no gravitational restoring torque:

$$\dot{\alpha}_5 = -(k_4 + k_5)\dot{\hat{x}}_5 + \hat{x}_5 \left(-k_4 k_5 - 1 + \frac{B_{2\phi}}{I_2} \right) + (k_4 k_5 + 1)\dot{x}_{4d} + (k_4 + k_5)\ddot{x}_{4d} + \dddot{x}_{4d}. \tag{37}$$

The yaw input gain vanishes at $\hat{x}_6 = -b_2/(2a_2) = -2.25$, again outside the physical envelope, and the same safeguard applies.

4.5. Controller Gains

The implemented gains are listed in Table 3.

Table 3. Implemented B-STA controller gains.

Parameter	Pitch Symbol	Value	Yaw Symbol	Value
Backstepping gain 1	k_1	8	k_4	20
Backstepping gain 2	k_2	6	k_5	20
Surface coefficient 1	c_1	5	c_4	5
Surface coefficient 2	c_2	3	c_5	3
STA proportional gain	λ_1	10	λ_3	10
STA integral gain	λ_2	5	λ_4	5
Anti-windup bound	$w_{\max}$	50	$w_{\max}$	50

The yaw backstepping gains are set substantially higher than the pitch gains. This reflects the structural asymmetry identified in Section 2: the horizontal channel has a moment of inertia roughly one third that of the vertical channel, a viscous friction coefficient roughly seventeen times larger, and no gravitational restoring torque to assist the return to a commanded attitude. Faster error convergence is therefore required on yaw to achieve comparable regulation, and the lower inertia permits it without excessive control effort. The same reasoning governs the gain search space of the scheduler in Section 6. Section 2.5

records the price of this choice: raising k_4 and k_5 also raises the gain at which the residual coupling reaches the yaw sliding variable.

5. Stability Analysis

The implemented closed loop is an interconnection of three subsystems: the observer error, the sliding variable, and the reduced tracking error. This section treats them as a cascade and establishes practical stability of the tracking error under bounded residual perturbation. Each element is stated separately so that the assumptions, the intermediate results, and the final conclusion can be examined independently.

5.1. Assumption on the Residual Perturbation

Assumption 1. *After application of (29), and while the command is not saturated, the sliding dynamics admit the perturbed super-twisting form*

$$\dot{s} = -\lambda_1 |s|^{1/2} \operatorname{sgn}(s) + w + \delta(t), \qquad \dot{w} = -\lambda_2 \operatorname{sgn}(s), \tag{38}$$

where the residual δ collects four contributions: the cross-coupling d_ψ discarded by the decentralised design and bounded in Section 2.5; the observer estimation error entering through the Lipschitz terms of φ_p and bounded through (12); parametric mismatch between Table 1 and the physical plant; and discretisation and quantisation effects. There exist constants $\Delta_0, \Delta_1 \geq 0$ with

$$|\delta(t)| \leq \Delta_0, \qquad |\dot{\delta}(t)| \leq \Delta_1 \tag{39}$$

almost everywhere on the compact operating set.

Remark 1. *Unlike the corresponding assumption in much of the sliding mode literature, these bounds are not postulated. Section 2.5 derives the coupling contribution from the mechanical stops, the actuator saturation, and the coefficients of Table 1, and evaluates it numerically for both axes; the observer contribution follows from (12). The derivative bound Δ_1 depends explicitly on the saturation limit $U_{\max}$, which is a structural property of the implementation. Deriving a bound rather than assuming one carries an obligation to report what the derivation yields, including where the resulting figure is unfavourable, and Section 5.2 does so.*

5.2. Convergence of the Sliding Variable

Lemma 1. *Consider (38) under Assumption 1 with $\Delta \geq \Delta_1$. If the super-twisting gains satisfy*

$$\lambda_2 > \Delta, \qquad \lambda_1 > \sqrt{\frac{2\Delta(\lambda_2 + \Delta)}{\lambda_2 - \Delta}}, \tag{40}$$

then the pair (s, w) converges in finite time to a compact neighbourhood of the origin whose radius is a class-$\mathcal{K}$ function of Δ_0 and Δ_1, collapsing to the origin when $\delta \equiv 0$.

Proof. Introduce $\boldsymbol{\xi} = \left[|s|^{1/2} \operatorname{sgn}(s),\ w\right]^\top$ and the strict Lyapunov function $V_{\mathrm{STA}} = \boldsymbol{\xi}^\top \mathbf{P} \boldsymbol{\xi}$ with $\mathbf{P} = \mathbf{P}^\top \succ 0$ satisfying the Lyapunov inequality associated with

$$\mathbf{A}_{\mathrm{STA}} = \begin{bmatrix} -\lambda_1 & 1 \\ -\lambda_2 & 0 \end{bmatrix}. \tag{41}$$

Under (40), the unperturbed dynamics give $\dot{V}_{\mathrm{STA}} \leq -\eta V_{\mathrm{STA}}^{1/2}$ for some $\eta > 0$, establishing finite-time convergence to the origin. Under the perturbation the same construction yields $\dot{V}_{\mathrm{STA}} \leq -\eta V_{\mathrm{STA}}^{1/2} + \rho(\Delta_0, \Delta_1)$, so convergence is to the ball on which the two terms

balance. The argument follows Moreno and Osorio [7]; see also [35] for the perturbed case. □

With the gains of Table 3, $\lambda_1 = 10$ and $\lambda_2 = 5$, the second inequality of (40) rearranges to $2\Delta^2 + (2\lambda_2 + \lambda_1^2)\Delta - \lambda_1^2\lambda_2 < 0$, whose positive root gives the admissible range

$$\Delta < 4.22. \tag{42}$$

Comparison with the evaluated bounds.

Table 2 gives $\Delta_1 = 105$ on pitch and $\Delta_1 = 627$ on yaw. Both exceed (42), and we state the consequence rather than adjust the gains retrospectively to meet a condition that the evaluated bounds do not support.

Three observations follow. First, (40) is sufficient and not necessary. Exceeding it withdraws the finite-time guarantee of Lemma 1 but does not imply divergence, and the sliding variables recorded in Section 8 remain bounded throughout the run. What Theorem 1 then asserts is boundedness of the tracking error through the input-to-state property of Lemma 2, which does not itself depend on (40), rather than convergence to a tight residual ball.

Second, the bounds of Table 2 are envelope quantities composed of simultaneous extremes. The dominant contribution to $|\dot{d}_\psi|$ is the term in $\dot{x}_5$, which is bounded at 22.2 rad/s^2 by the saturation limit but reaches roughly 0.25 rad/s^2 on the sinusoidal references actually commanded. The operative perturbation on the recorded runs is therefore smaller than the tabulated figure by roughly two orders of magnitude on that term, which is consistent with the observed behaviour and is why the design functions in practice.

Third, and independently of the conservatism, the *ratio* between the axes is informative because the same bounding procedure is applied to both. The prediction is that the yaw sliding variable operates further outside the guaranteed regime than the pitch sliding variable by a factor of approximately six, and this is precisely what the data show. In simulation the pitch surface converges to a tight neighbourhood of zero within approximately three seconds and remains there, while the yaw surface retains a persistent oscillation of order ± 20 rad/s^2 for the entire run, as reported in Section 7; the same ordering is recorded on hardware in Section 8. The evaluated bounds therefore account for a qualitative difference between the two axes that would otherwise have to be attributed to unmodelled effects alone, and they identify the mechanism as the reaction-torque coefficient $1.75|K_c|/I_2$ acting through the elevated yaw transfer gain. Satisfying (40) on the yaw axis would require either a substantial increase in λ_3 and λ_4 or a reduction of $(c_5 + k_4 + k_5)$, the second of which conflicts with the tracking requirement that motivated the high yaw gains. This trade-off is identified in Section 9 as the principal open item for the yaw channel.

5.3. Input-to-State Stability of the Reduced Error Dynamics

Lemma 2. *The reduced tracking dynamics (25) are input-to-state stable with respect to s. Specifically,* $\mathbf{A}_{\mathrm{cl}}$ *is Hurwitz for all positive* k_1, k_2, c_1, c_2*, and for the symmetric* $\mathbf{P}_e \succ 0$ *solving* $\mathbf{A}_{\mathrm{cl}}^\top \mathbf{P}_e + \mathbf{P}_e \mathbf{A}_{\mathrm{cl}} = -\mathbf{Q}$ *with* $\mathbf{Q} = \mathbf{Q}^\top \succ 0$*, the tracking error satisfies*

$$\|\mathbf{e}_{12}(t)\| \le \beta\big(\|\mathbf{e}_{12}(0)\|, t\big) + \gamma \sup_{\tau \le t} |s(\tau)|, \qquad \gamma = \frac{2\lambda_{\max}(\mathbf{P}_e)}{\lambda_{\min}(\mathbf{Q})} \sqrt{\frac{\lambda_{\max}(\mathbf{P}_e)}{\lambda_{\min}(\mathbf{P}_e)}}, \tag{43}$$

where $\mathbf{e}_{12} = [e_1, e_2]^\top$ *and* β *is of class* $\mathcal{KL}$*.*

Proof. The characteristic polynomial of $\mathbf{A}_{\mathrm{cl}}$ is

$$p(\lambda) = \lambda^2 + (k_1 + k_2 + c_2)\lambda + (k_1 k_2 + k_1 c_2 + 1 + c_1), \tag{44}$$

whose coefficients are strictly positive for positive gains, so $\mathbf{A}_{\mathrm{cl}}$ is Hurwitz by the Routh criterion. Taking $V_e = \mathbf{e}_{12}^{\top}\mathbf{P}_e\mathbf{e}_{12}$ and differentiating along (25),

$$\dot{V}_e = -\mathbf{e}_{12}^{\top}\mathbf{Q}\mathbf{e}_{12} + 2\mathbf{e}_{12}^{\top}\mathbf{P}_e[0,\ 1]^{\top}s \le -\lambda_{\min}(\mathbf{Q})\|\mathbf{e}_{12}\|^2 + 2\lambda_{\max}(\mathbf{P}_e)\|\mathbf{e}_{12}\||s|, \tag{45}$$

which is negative whenever $\|\mathbf{e}_{12}\| > 2\lambda_{\max}(\mathbf{P}_e)|s|/\lambda_{\min}(\mathbf{Q})$. Standard comparison arguments [4] give (43). □

For the pitch gains of Table 3, (44) gives $p(\lambda) = \lambda^2 + 17\lambda + 78$ with eigenvalues $\lambda_{1,2} = -8.5 \pm j2.398$. For the yaw gains it gives $p(\lambda) = \lambda^2 + 43\lambda + 466$ with eigenvalues $\lambda_{1,2} = -21.5 \pm j1.936$. Both pairs are Hurwitz, with damping ratios of 0.963 and 0.996 respectively.

Lemma 2 holds for all t, not only after the sliding set is reached. This is the reason for retaining the exact relation (25) rather than the algebraic constraint $e_3 = -c_1 e_1 - c_2 e_2$ valid only on $s = 0$: the reaching phase is covered by the same argument. It is also the reason the tracking result survives the comparison of Section 5.2, since a bounded sliding variable suffices to bound the tracking error even when the residual ball is not tight.

5.4. Cascade Result

Theorem 1. *Consider the plant (2) under the observer (10) and the control law (29) with positive gains satisfying (40), and let Assumption 1 hold. Then,*

(i) The estimation error $\tilde{\mathbf{x}}$ converges exponentially to a residual set of radius κ_o, by (12).
(ii) The sliding pair (s, w) converges in finite time to a ball of radius $r(\Delta_0, \Delta_1)$, by Lemma 1.
(iii) The tracking error satisfies the ultimate bound

$$\limsup_{t\to\infty} \|\mathbf{e}_{12}(t)\| \le \gamma\, r(\Delta_0, \Delta_1), \tag{46}$$

with γ given in (43), by Lemma 2.
(iv) The unsaturated command $U_1(t)$ is a continuous function of time, because the discontinuity appears in $\dot{w}$ and not in U_1.

In the ideal case $\delta \equiv 0$ and $\kappa_o = 0$, the residual sets collapse and asymptotic tracking is recovered.

Proof. Part (i) is the observer property. Part (ii) is Lemma 1. For part (iii), the interconnection is a cascade: the sliding subsystem is driven by δ, which by Assumption 1 incorporates the observer error, and the reduced subsystem is driven by s through Lemma 2. Since an input-to-state stable system driven by a signal converging to a ball has state converging to a ball scaled by the ISS gain [4], and since the sliding subsystem converges in finite time, (46) follows. Part (iv) is immediate from (29) and (30) together with the continuity of $|s|^{1/2}\operatorname{sgn}(s)$ at the origin. □

Remark 2. *Theorem 1 establishes practical stability, not global asymptotic stability of the exact origin. This is a deliberate weakening relative to the claim made in our earlier work on this platform [29], and it reflects what can honestly be asserted for a decentralised design operating on a coupled physical plant through estimated states. The result is stated in terms of an explicit ultimate bound so that the contribution of each residual term can be traced.*

Remark 3. *Part (ii), and through it the tightness of the bound in part (iii), requires (40). Section 5.2 records that the evaluated Δ_1 exceeds the admissible range on both axes, so on the reported runs the conclusion that is supported is boundedness of (s, w) and hence of $\mathbf{e}_{12}$ through Lemma 2, rather than finite-time arrival at the tight residual ball. We separate the two so that the experimental results of Section 8 are read against the claim the analysis actually supports.*

Remark 4. *Part (iv) refers to the unsaturated command. During intervals in which the saturation in (29) is active the applied voltage is clipped and the continuity property does not extend to the applied signal. Section 8 reports the fraction of the experimental run spent at the saturation limit so that the practical extent of this qualification can be assessed.*

6. Windowed Grey Wolf Gain Scheduling

6.1. Motivation and Architecture

The B-STA controller of Section 4 uses fixed gains chosen once and held for the duration of a run. An alternative is to adapt the backstepping gains online while the plant is running. This section develops such a scheme, based on the Grey Wolf Optimiser [26], and establishes a structural condition that any online metaheuristic gain scheduler must satisfy in order to function at all.

The architecture has two layers. The inner layer is a decentralised three-step backstepping law without sliding mode injection, parameterised by the gain vector $\boldsymbol{\theta} = [k_1, k_2, k_3]^\top$ on pitch and $[k_4, k_5, k_6]^\top$ on yaw. With the gains available at sample k, the pitch voltage is

$$U_1 = \frac{1}{g(\hat{\mathbf{x}})}\left[\dot{\alpha}_2 - \dot{\mathcal{E}}_{\text{nom}} - k_3 e_3 - e_2\right]. \tag{47}$$

The outer layer is a Grey Wolf Optimiser that adjusts $\boldsymbol{\theta}$ during operation. Two independent instances run in parallel, one per axis, each minimising its own subsystem error. Robustness is pursued through gain adaptation rather than through discontinuous injection, so no sliding surface appears in (47). The symbols k_1 to k_5 denote the same gain positions within the same recursive chains as in Table 3, so that the fixed-gain and scheduled designs remain directly comparable; the values they take differ between the two designs, and k_3 and k_6 have no counterpart in B-STA, where the third recursive step is replaced by the composite surface.

6.2. Why the Fitness Must Be Evaluated over a Window

Before specifying the scheduler we establish a condition that constrains its design. The result is elementary, but it has a decisive practical consequence and appears not to be stated explicitly in the online metaheuristic tuning literature.

Proposition 1. *Consider a population-based online tuner that, at sample k, scores each candidate $\boldsymbol{\theta}_i$, $i = 1, \ldots, N$, by a fitness of the form*

$$f(\boldsymbol{\theta}_i) = g\big(e(k)\big) + h(\boldsymbol{\theta}_i), \tag{48}$$

where $e(k)$ is the tracking error measured at sample k and the candidates have not yet been applied to the plant. Then the ranking induced by f over the population is identical to the ranking induced by h alone, and the tuner converges to $\arg\min_{\boldsymbol{\theta}} h(\boldsymbol{\theta})$ irrespective of the closed-loop behaviour of the plant.

Proof. At sample k the quantity $e(k)$ is a single realised scalar, determined by the trajectory up to that instant and by the gains already applied. It does not depend on the index i. Hence $g(e(k))$ is a constant offset common to every member of the population, and for

any pair i, j we have $f(\boldsymbol{\theta}_i) < f(\boldsymbol{\theta}_j)$ if and only if $h(\boldsymbol{\theta}_i) < h(\boldsymbol{\theta}_j)$. The selection operator of any rank-based metaheuristic is therefore invariant to $g(e(k))$, and the search minimises h alone. □

The practical reading of Proposition 1 is that a regularisation term intended as a mild secondary penalty becomes the sole objective. If h penalises deviation from a nominal gain vector $\bar{\boldsymbol{\theta}}$, the scheduler converges to $\bar{\boldsymbol{\theta}}$ and remains there, producing a controller that is nominally adaptive and effectively fixed.

The remedy is structural rather than a matter of gain choice. The fitness must be made candidate-dependent, and on a physical plant the only way to achieve this is to apply each candidate and measure the error that results. This requires an evaluation window, which in turn requires the sequential test-and-settle mechanism developed below. Windowing is therefore not a computational convenience introduced to spread the population evaluation over time; it is what makes the optimisation well-posed.

6.3. Windowed Sequential Evaluation

Optimisation proceeds by alternating two phases per candidate.

- **Test phase.** Candidate $\boldsymbol{\theta}_i$ is applied to the plant for T_w samples, during which the integral of time-weighted absolute error is accumulated for that candidate alone.
- **Settle phase.** The current best gain vector $\boldsymbol{\alpha}$ is restored for an equal window, so that the next candidate begins from a comparable plant state and any transient excited by the previous candidate has decayed. This phase also serves an experimental safety function.

The fitness of candidate i, evaluated at the close of its test window beginning at sample k_0, is

$$J(\boldsymbol{\theta}_i) = \underbrace{\sum_{n=k_0}^{k_0+T_w-1} (nT_s)\,|e(n)|\,T_s}_{\text{windowed ITAE, candidate-dependent}} + \beta \sum_{j=1}^{3} \left(\frac{\theta_{i,j} - \bar{\theta}_j}{\Delta\theta_j} \right)^2. \tag{49}$$

The first term now depends on i, because the error samples $e(n)$ over the window are those produced by candidate i acting on the plant. Proposition 1 therefore no longer applies, and the regularisation term resumes its intended secondary role of discouraging boundary gains.

Applied gains pass through a first-order rate limiter

$$\boldsymbol{\theta}_{\text{applied}}(k) = \boldsymbol{\theta}_{\text{applied}}(k-1) + \lambda \Big[\boldsymbol{\theta}_{\text{target}}(k) - \boldsymbol{\theta}_{\text{applied}}(k-1) \Big], \qquad \lambda = 0.1, \tag{50}$$

where $\boldsymbol{\theta}_{\text{target}}$ is the candidate under test during a test phase and $\boldsymbol{\alpha}$ during a settle phase. With $\lambda = 0.1$ the discrete time constant is approximately ten samples, corresponding to 0.1 s at the control period of Section 8, so the applied gains reach the target well within a window of $T_w = 50$ samples while the transition remains smooth enough to avoid actuator jerk.

6.4. Grey Wolf Update

When all N candidates of a generation have been scored, the population is ranked and the three best members are designated $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ and $\boldsymbol{\delta}$. Each agent is repositioned as the mean of three leader-guided estimates,

$$\boldsymbol{\theta}_i^{(g+1)} = \tfrac{1}{3}(\mathbf{X}_1 + \mathbf{X}_2 + \mathbf{X}_3), \qquad \mathbf{X}_m = \boldsymbol{\ell}_m - \mathbf{A}_m \odot \left| \mathbf{C}_m \odot \boldsymbol{\ell}_m - \boldsymbol{\theta}_i^{(g)} \right|, \tag{51}$$

with $\ell_1 = \alpha$, $\ell_2 = \beta$, $\ell_3 = \delta$, and

$$\mathbf{A}_m = 2a\,\mathbf{r}_1 - a, \qquad \mathbf{C}_m = 2\mathbf{r}_2, \qquad a = 2\left(1 - \frac{g}{G_{\max}}\right), \tag{52}$$

where $\mathbf{r}_1$ and $\mathbf{r}_2$ are drawn uniformly from $[0,1]^3$. Early generations with $a > 1$ favour exploration, and later generations local refinement. Positions are projected onto the box $[\mathbf{lb}, \mathbf{ub}]$ after each update. The complete procedure is given in Algorithm 1.

Algorithm 1 Windowed GWO gain scheduling, executed per subsystem

Require: Error $e(k)$, bounds $[\mathbf{lb}, \mathbf{ub}]$, window T_w, generations $G_{\max}$
Ensure: Rate-limited gains $\boldsymbol{\theta}_{\text{applied}}(k)$
Initialisation, executed once
Initialise N agents uniformly in $[\mathbf{lb}, \mathbf{ub}]$
$\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\delta}, \boldsymbol{\theta}_{\text{applied}} \leftarrow \bar{\boldsymbol{\theta}}$ ▷ safe midrange start
evalIdx $\leftarrow 1$; evalMode $\leftarrow$ TEST; winCtr $\leftarrow 0$; $J_{\text{acc}} \leftarrow 0$; $g \leftarrow 1$
Per-sample execution
winCtr $\leftarrow$ winCtr $+ 1$
if evalMode $=$ TEST **then**
    $\boldsymbol{\theta}_{\text{target}} \leftarrow \boldsymbol{\theta}_{\text{evalIdx}}$
    $J_{\text{acc}} \leftarrow J_{\text{acc}} + (kT_s)|e(k)|T_s$ ▷ candidate-specific ITAE
else
    $\boldsymbol{\theta}_{\text{target}} \leftarrow \boldsymbol{\alpha}$
end if
$\boldsymbol{\theta}_{\text{applied}} \leftarrow \boldsymbol{\theta}_{\text{applied}} + \lambda(\boldsymbol{\theta}_{\text{target}} - \boldsymbol{\theta}_{\text{applied}})$
if winCtr $\geq T_w$ **then**
    winCtr $\leftarrow 0$
    if evalMode $=$ TEST **then**
        Fit(evalIdx) $\leftarrow J_{\text{acc}} + \beta \sum_j \left((\theta_{\text{evalIdx},j} - \bar{\theta}_j)/\Delta\theta_j\right)^2$
        $J_{\text{acc}} \leftarrow 0$; evalMode $\leftarrow$ SETTLE
    else
        evalMode $\leftarrow$ TEST; evalIdx $\leftarrow$ evalIdx $+ 1$
        if evalIdx $> N$ **then**
            Rank population; update $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\delta}$ by cascading comparison
            $a \leftarrow \max\left(0,\ 2(1 - g/G_{\max})\right)$
            for $i = 1$ **to** N **do**
                Update $\boldsymbol{\theta}_i$ via (51) and (52); project onto $[\mathbf{lb}, \mathbf{ub}]$
            end for
            Fit $\leftarrow \infty$; evalIdx $\leftarrow 1$; $g \leftarrow \min(g + 1, G_{\max})$
        end if
    end if
end if
return $\boldsymbol{\theta}_{\text{applied}}$

6.5. Search Space, Timing Budget and Safety

The two subsystems use different gain boxes, for the reason set out in Section 4: the horizontal channel has roughly one third the inertia, roughly seventeen times the viscous friction, and no gravitational restoring torque, so it requires substantially larger gains. The fixed-gain B-STA design uses $k_4 = k_5 = 20$ on that axis, which lies outside the pitch box entirely. Confining the yaw scheduler to the pitch box would therefore constrain it to a quarter of the loop gain that the same channel demonstrably requires. The boxes are listed in Table 4, and the regularisation centre and width follow each box so that the penalty remains a mild secondary term rather than a bias toward a fixed point.

Table 4. Windowed GWO configuration.

Parameter	Pitch	Yaw	Description
Population size N	10	10	Agents per axis
Maximum generations $G_{\max}$	20	20	Exploration horizon
Evaluation window T_w	50	50	Window, samples
Rate-limiter coefficient λ	0.1	0.1	Gain smoothing
Lower bound **lb**	$[0.1, 0.1, 0.1]$	$[0.1, 0.1, 0.1]$	Box floor
Upper bound **ub**	$[5, 2, 1]$	$[20, 20, 5]$	Box ceiling
Regularisation centre $\bar{\boldsymbol{\theta}}$	$[2.55, 1.05, 0.55]$	$[10.05, 10.05, 2.55]$	Box midpoint
Regularisation width $\Delta\boldsymbol{\theta}$	$[4.9, 1.9, 0.9]$	$[19.9, 19.9, 4.9]$	Box width
Regularisation weight β	0.01	0.01	Penalty weight

Each generation consumes $N \times 2 \times T_w$ samples, which with the configuration of Table 4 is $10 \times 2 \times 50 = 1000$ samples. At the control period $T_s = 0.01$ s of Section 8 this is 10 s of plant time per generation. The scheduler is therefore a slow outer loop: it advances one sample of the evaluation state machine at each control tick and updates the population only at generation boundaries. It is not a complete metaheuristic search executed at every sample.

This timing budget has a consequence that must be stated, because it governs how the recorded gain trajectories are to be read. A 100 s run accommodates ten generations, whereas the exploration schedule in (52) is written for $G_{\max} = 20$. The coefficient a therefore decays only from 1.90 to 1.00 across the run and never enters the regime $a < 1$ in which $|\mathbf{A}_m| < 1$ for all draws and the update contracts toward the leaders. The search remains in its exploratory phase for the entire experiment. This accounts directly for the sustained gain excursions reported in Section 7, which are the intended behaviour of the algorithm in that phase rather than an instability, and it identifies a concrete configuration error: with a 100 s horizon, $G_{\max}$ should be set to the number of generations the run can actually complete, or the horizon extended, if the refinement phase is to be reached at all. Section 8.5.4 returns to this point in interpreting the measured B-GWO performance.

The applied gains are rate-limited by (50), and a supervisory freeze to the incumbent $\boldsymbol{\alpha}$ is triggered if a candidate drives the measured error outside a preset envelope during its test window. These provisions bound the excursions a candidate can produce but do not constitute a closed-loop stability certificate for the time-varying-gain plant, and none is claimed. A complete stability proof for online metaheuristic scheduling on a nonlinear plant remains an open problem.

7. Numerical Validation

7.1. Configuration

Both architectures were exercised in MATLAB and Simulink R2014a before deployment. The plant used for simulation is the fully coupled model (2), including both gyroscopic terms, so that the simulation exercises exactly the coupling the decentralised controllers do not cancel. Integration is by classical fourth-order Runge–Kutta with zero-order-hold control at the implementation sample period $T_s = 0.01$ s, over a 100 s horizon, with voltages saturated at ± 2.5 V as on the hardware. Both axes track sinusoidal references of period 12.57 s with amplitudes 0.40 rad on pitch and 1.00 rad on yaw, identical for all three controllers.

The simulation omits mechanical cable tension, bearing stiction, and aerodynamic turbulence, and is used for two purposes only: verification of nominal behaviour, and controlled disturbance rejection testing under perturbations that cannot be applied repeatably to the physical rig.

7.2. Nominal Tracking

Figures 3 and 4 show nominal B-STA tracking and the evolution of the two sliding surfaces. Figure 5 shows the corresponding B-GWO response, and Figure 6 the gain trajectories produced by the windowed scheduler of Algorithm 1.

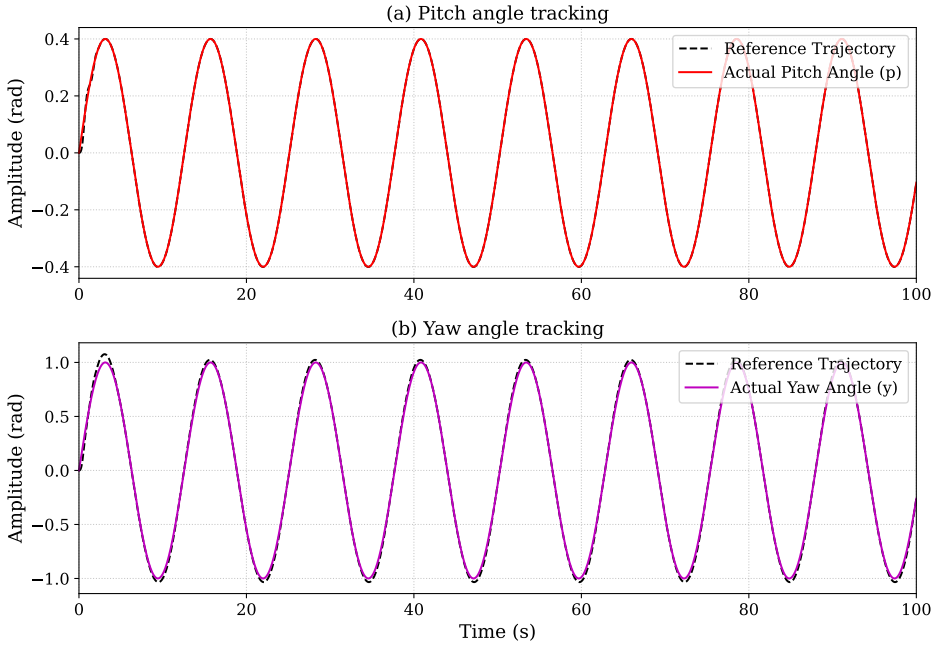


Figure 3. Nominal B-STA tracking on the fully coupled model. (**a**) Pitch angle ψ (rad) against reference ψ_d. (**b**) Yaw angle ϕ (rad) against reference ϕ_d.

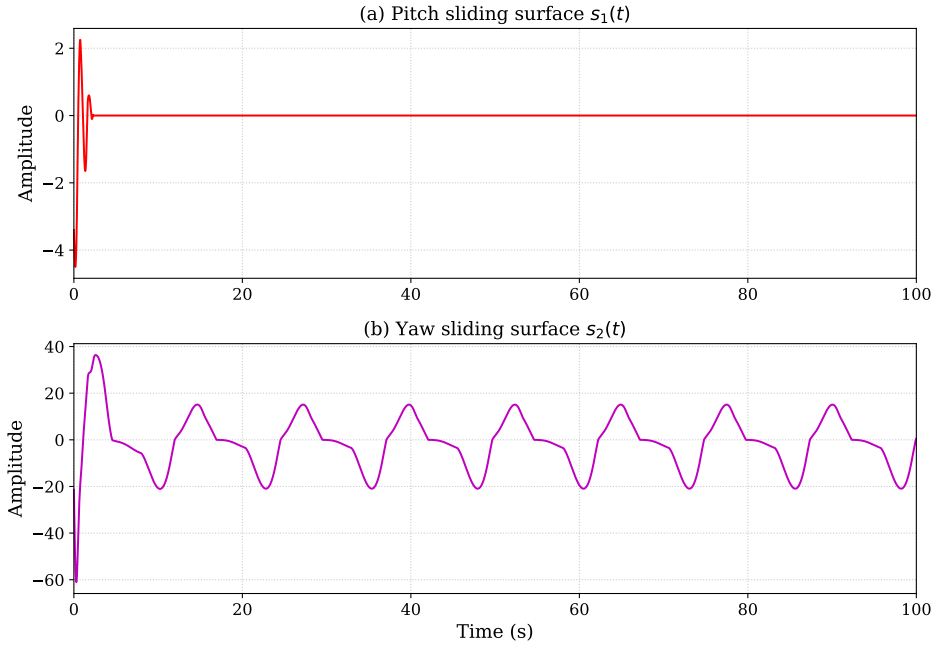


Figure 4. Nominal evolution of the B-STA sliding surfaces. (**a**) Pitch surface s (rad/s^2), which converges to a tight neighbourhood of zero within approximately three seconds. (**b**) Yaw surface s_{yaw} (rad/s^2), which remains bounded but retains a persistent oscillation for the duration of the run. The contrast is consistent with the residual bounds evaluated in Table 2 and discussed in Section 5.2.

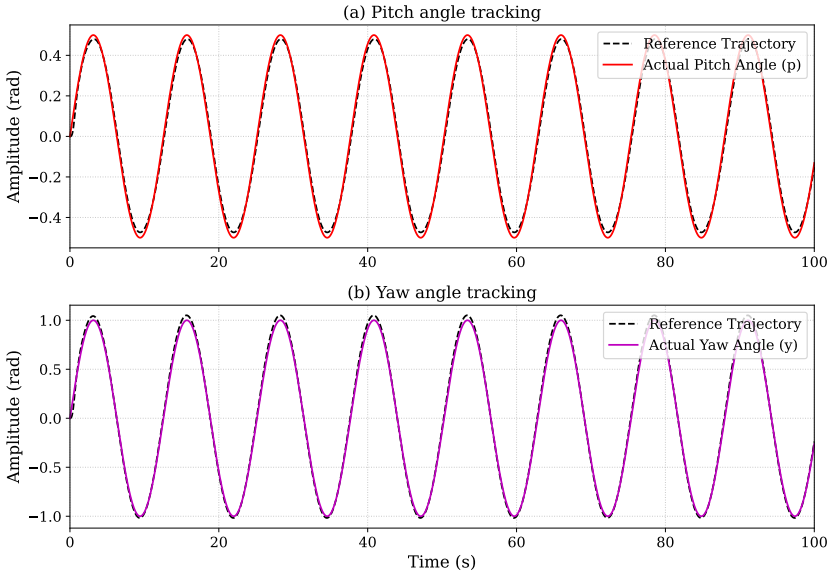


Figure 5. Nominal B-GWO tracking on the fully coupled model. (**a**) Pitch angle ψ (rad). (**b**) Yaw angle ϕ (rad).

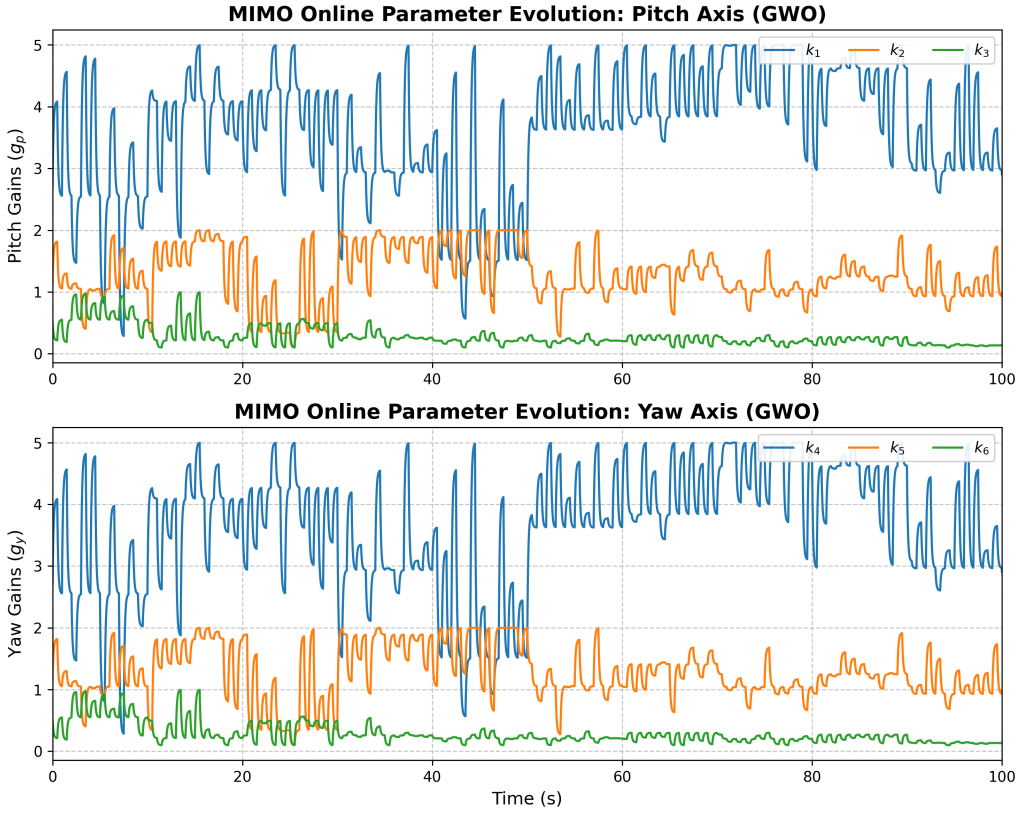


Figure 6. Gain trajectories produced by the windowed GWO scheduler under nominal conditions, one instance per axis. (**Upper panel**) pitch axis, gains $[k_1, k_2, k_3]$. (**Lower panel**) yaw axis, gains $[k_4, k_5, k_6]$. The rapid excursions correspond to successive test windows, in which individual candidates are applied to the plant and scored; the intervening returns toward a common level correspond to settle windows at the incumbent best. The excursions persist for the whole run because, as shown in Section 6.5, a 100 s horizon completes only ten of the twenty generations for which the exploration schedule was written, so the search does not reach its refinement phase. The two instances are independent: each ranks its own population against its own subsystem error, and neither leader set influences the other.

7.3. Disturbance Rejection

Robustness is the property the continuous second-order injection is intended to provide, and nominal tracking does not test it. Repeatable disturbance injection on the physical rig would require instrumentation not available on this platform, so three perturbation scenarios were applied in simulation on the fully coupled model. Each is applied identically to B-STA, to B-GWO, and to the first-order Backstepping–Sliding Mode controller of [29], so that the comparison isolates the effect of the injection structure.

- **Scenario D1, matched load step.** An additive disturbance of 0.5 V is injected at the pitch actuator input at $t = 50$ s, after the saturation stage so that it acts as a genuine actuator-side perturbation, and held for the remainder of the run. This emulates a change in beam loading or a shift in the counterweight and enters through the same channel as the control, so it directly tests the compensating action of the super-twisting integrator state $w(t)$.
- **Scenario D2, parametric perturbation.** The plant moments of inertia I_1 and I_2 are increased by 20% and the aerodynamic coefficients a_1 and a_2 reduced by 15%, while all three controllers retain the nominal values of Table 1. The mismatch is therefore genuine and uncompensated.
- **Scenario D3, aerodynamic turbulence.** A band-limited random torque of standard deviation equal to 10% of the nominal thrust term, with bandwidth 0 to 5 Hz, is added to both angular acceleration equations. A fixed seed is used so that every controller experiences the same realisation.

Table 5 reports the resulting tracking error. For D1 the figures are post-event root mean square errors over $t \in [50, 100]$ s; for D2 and D3 they are full-run values. Figure 7 shows the response to Scenario D1.

Table 5. Disturbance rejection on the fully coupled simulation model. Post-event RMSE is computed over $t \in [50, 100]$ s for D1 and over the full run for D2 and D3. Recovery time is the interval after the D1 event before the pitch error returns within 1.5 times its pre-event root mean square value and remains there. Bold denotes the best value in each column within each scenario; the three controllers track identical references, so the comparison is direct.

Scenario	Controller	Pitch RMSE (rad)	Yaw RMSE (rad)	Recovery (s)
D1, matched load step	BSMC	0.0079	**0.0310**	6.14
	B-GWO	0.2876	0.3174	not recovered
	B-STA	**0.0002**	0.0750	**2.62**
D2, parametric	BSMC	0.0187	**0.0365**	n/a
	B-GWO	0.0501	0.1686	n/a
	B-STA	**0.0102**	0.0849	n/a
D3, turbulence	BSMC	0.0440	**0.0535**	n/a
	B-GWO	0.0670	0.2117	n/a
	B-STA	**0.0344**	0.0957	n/a

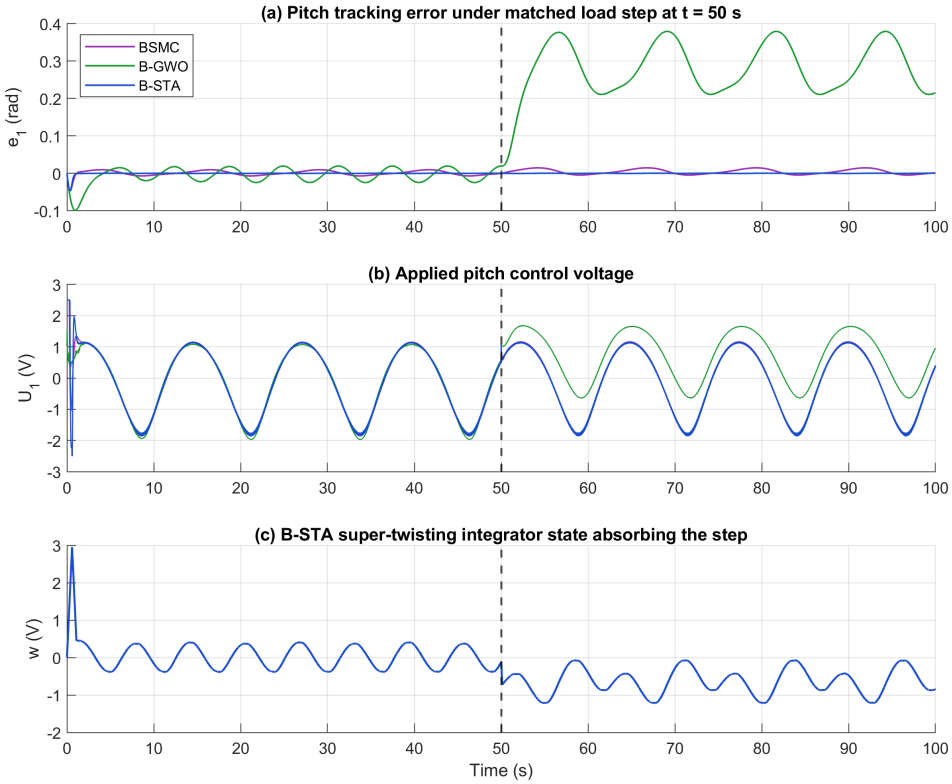


Figure 7. Response to Scenario D1, a matched load step of 0.5 V applied at the pitch actuator at $t = 50$ s. (**a**) Pitch tracking error e_1 (rad). (**b**) Applied pitch control voltage U_1 (V). (**c**) B-STA super-twisting integrator state w (V), which migrates from a mean of $+0.006$ V before the event to -0.643 V after it, offsetting the injected disturbance.

7.4. Discussion of the Disturbance Results

Three observations follow from Table 5 and Figure 7.

The compensation mechanism is confirmed directly. Under the matched step, the super-twisting integrator state migrates from a mean of $+0.006$ V before the event to -0.643 V after it, a shift in the direction required to oppose a positive actuator disturbance and of the magnitude expected once scaled by $g^{-1}(\hat{\mathbf{x}})$. The integration in (30) continues until the sliding variable returns to zero, so the offset is cancelled exactly rather than merely attenuated. The mean residual pitch error after the event is -2×10^{-5} rad, and the post-event RMSE of 0.0002 rad is the smallest in the table by more than an order of magnitude. Because the discontinuity acts on $\dot{w}$, this compensation is achieved without any increase in the high-frequency content of the applied voltage, which is visible in panel (b).

B-GWO provides the controlled comparison that isolates the mechanism. The B-GWO inner law (47) is the same recursive backstepping structure as B-STA but without the second-order injection. Under the matched step it retains a permanent pitch offset of 0.282 rad and never recovers, which is the expected behaviour of a smooth law with no integral action against a persistent matched disturbance. The contrast with B-STA, which shares the same recursive structure and differs only in the final step, attributes the rejection specifically to the super-twisting term rather than to the backstepping design or to the gain values.

On the yaw column, the baseline is the better performer in all three scenarios. This is the opposite of the ordering observed on hardware in Section 8, where B-STA yaw RMSE of 0.1043 rad improved on the baseline value of 0.1858 rad. We report the result as obtained and offer the following interpretation rather than the reverse. Two factors act in the same direction. The first is analytical: Section 5.2 shows that the residual perturbation entering the yaw sliding variable exceeds the range admitted by (40) by a wider margin than on

pitch, by a factor of approximately six in Δ_1, so the yaw surface is not expected to reach the tight residual ball that the pitch surface reaches, and Figure 4 confirms that it does not. The second is that the simulation omits stiction, cable tension and the friction nonlinearities that dominate the horizontal channel on the physical rig, and these are precisely the effects a first-order boundary-layer injection handles poorly. Removing them removes the conditions under which the second-order injection is advantageous on that axis. The pitch column, where the gravitational load is retained in the model and is the dominant term and where the residual bound is smallest, shows B-STA ahead in all three scenarios.

Two limitations of this study should be recorded. The B-GWO rows are single realisations of a stochastic search and vary between runs; multi-trial characterisation is identified as future work in Section 9. And the total variation in the simulated B-STA command is not reported here, because at a 10 ms step with a sliding variable very close to zero the numerical toggling of $\mathrm{sgn}(s)$ dominates that statistic without any corresponding excursion in the applied voltage, as panel (b) of Figure 7 shows. Actuator load is therefore quantified on hardware in Section 8, where it is physically meaningful.

8. Real-Time Experimental Validation

8.1. Experimental Configuration

The experiments were conducted on the physical two-degree-of-freedom TRMS. The control algorithms and the observer of Section 3 execute under the Simulink Desktop Real-Time kernel at the control period $T_s = 0.01$ s, corresponding to a control rate of 100 Hz. An Advantech PCI-1711 board provides twelve-bit digital-to-analogue conversion of the motor voltages and reads the two incremental optical encoders through its counter inputs. All commands are saturated at ± 2.5 V in software before conversion. Figure 8 shows the signal path.

The recorded B-GWO run was produced by the windowed implementation of Algorithm 1 with the configuration of Table 4. The single-sample fitness variant described in Section 8.5.4 was executed as a separate diagnostic and is reported there for its bearing on Proposition 1; it does not contribute to any of the results reported below.

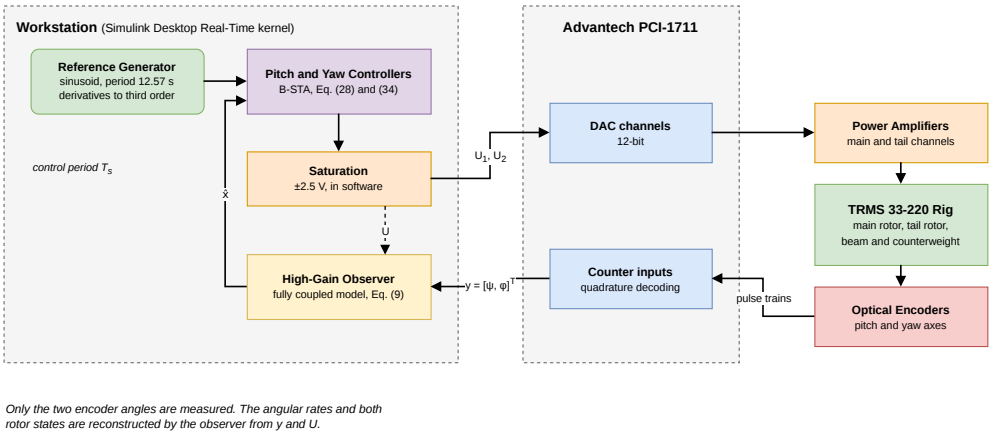


Figure 8. Real-time experimental architecture. The workstation executes the observer and both controllers under the Simulink Desktop Real-Time kernel. Control voltages pass through the digital-to-analogue channels of the Advantech PCI-1711 board to the motor power amplifiers, and the two encoder channels return through the board counter inputs. Saturation at ± 2.5 V is applied in software before conversion.

8.2. Reference Trajectories

Both axes track sinusoidal references of period 12.57 s, corresponding to 0.0796 Hz. The amplitudes differ between the two recorded runs and are stated in Table 6. Because tracking error scales with commanded excursion, all cross-controller comparisons in Section 8.5 report normalised error alongside absolute error.

Table 6. Reference trajectory amplitudes for the recorded experimental runs.

Run	Pitch Amplitude (rad)	Yaw Amplitude (rad)	Period (s)
B-STA	0.40	1.00	12.57
B-GWO	0.50	0.60	12.57

8.3. Performance Metrics and Reporting Protocol

Four indices are computed over the 100 s recording window:

- **Root mean square error**, $\mathrm{RMSE} = \left(N^{-1}\sum_k e^2(kT_s)\right)^{1/2}$, together with the normalised value RMSE/A where A is the reference amplitude of that axis.
- **Integral of time-weighted absolute error**, $\mathrm{ITAE} = \sum_k (kT_s)|e(kT_s)|T_s$.
- **Control effort**, $\mathrm{CE} = \sum_k U^2(kT_s)T_s$, reported with the sampling-rate-invariant mean square voltage $\overline{U^2} = \mathrm{CE}/T$.
- **Total variation in the applied voltage**, $\mathrm{TV} = \sum_k |U(k+1) - U(k)|$, a direct measure of cumulative actuator excursion, reported with the fraction of the run spent at the saturation limit.

Two points of the protocol require explicit statement.

First, all sums are evaluated with the control period $T_s = 0.01$ s read from the recorded time vector rather than assumed. During preparation of this revision we identified that the real-time ITAE values reported in our earlier work on this platform [29] were computed with a sampling interval that did not match the rate at which the controller output is updated, which inflated them by a factor of ten. Root mean square error, being an average rather than a sum, was unaffected, which is why the discrepancy was not apparent from that table alone. The baseline ITAE values in Table 7 are recomputed from the raw logs and supersede the values printed in [29]. A corrigendum to that article has been notified to the Editorial Office. As a consistency check, every entry in Table 7 satisfies the Cauchy–Schwarz bound

$$\mathrm{ITAE} \leq \frac{T^2}{\sqrt{3}}\,\mathrm{RMSE}, \tag{53}$$

which any consistently computed pair must obey and which we recommend as a routine verification for time-weighted error reporting.

Second, all quantitative results correspond to a single experimental run per controller. No multi-trial mean or standard deviation is reported. The values are comparative evidence under matched laboratory conditions rather than a statistical population study, and the stochastic character of the metaheuristic scheduler in particular warrants multi-trial characterisation in future work.

8.4. Results

Figures 9–12 show the recorded responses, and Table 7 consolidates the metrics.

Table 7. Real-time performance over a single 100 s window. Normalised RMSE is referred to the reference amplitude of that axis in Table 6. The two rows taken from [29] were recorded under a different experimental campaign whose commanded amplitudes are not stated in that article, so no normalised value is entered for them and the comparison with those rows is made on absolute error only. Baseline ITAE values are recomputed from the raw logs of [29] and supersede the values printed there. Bold denotes the best value in each column on each axis, applied mechanically. The absolute RMSE, ITAE, CE and TV columns are not amplitude-corrected and the two runs of the present study commanded the different excursions listed in Table 6; Section 8.5.1 sets out what that does and does not permit to be concluded.

Controller	Axis	RMSE (rad)	RMSE/A (%)	ITAE (rad·s)	CE (V^2·s)	TV (V)
Hybrid BSMC [29]	Pitch	0.0682	n/a	273.37	390.47	2416.7
	Yaw	0.1858	n/a	862.47	95.48	**154.3**
Online RS-QBHO [29]	Pitch	0.0799	n/a	338.23	n/a	n/a
	Yaw	0.2305	n/a	1026.86	n/a	n/a
Backstepping–GWO	Pitch	0.0858	17.16	354.71	207.57	257.7
	Yaw	**0.0683**	11.39	**274.17**	110.08	714.8
Backstepping–STA	Pitch	**0.0318**	**7.94**	**108.04**	**105.94**	**82.4**
	Yaw	0.1043	**10.43**	405.18	**48.21**	162.7

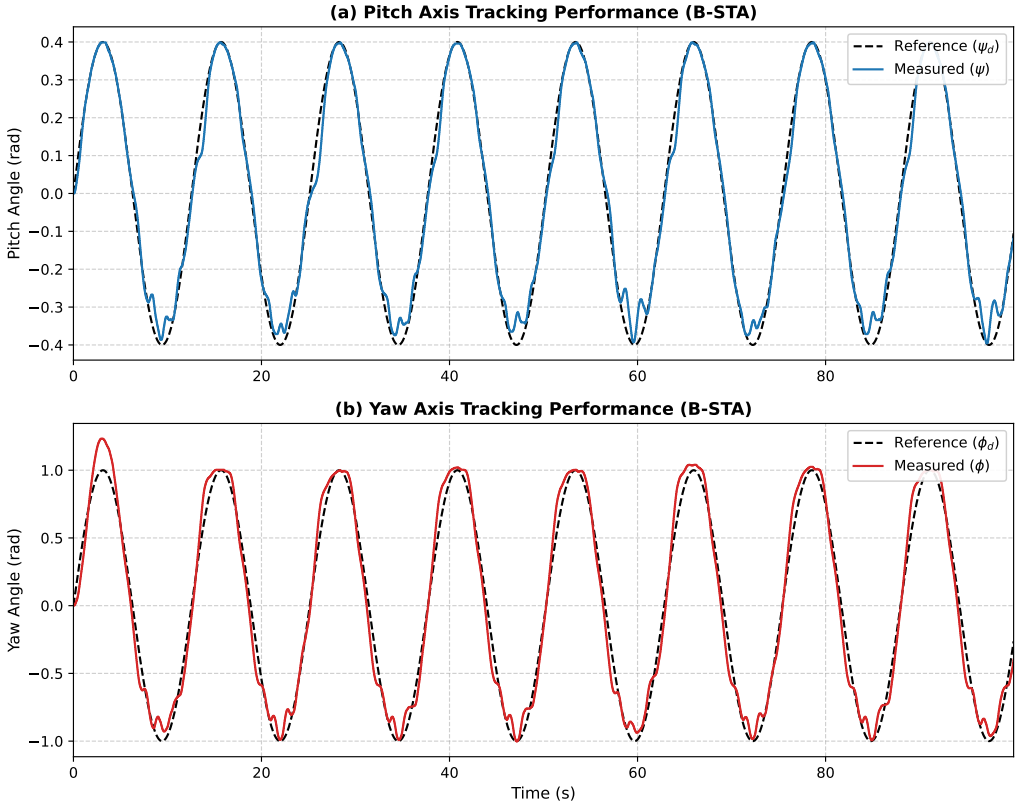


Figure 9. Real-time B-STA tracking. (**a**) Pitch angle ψ (rad) against reference ψ_d. (**b**) Yaw angle ϕ (rad) against reference ϕ_d.

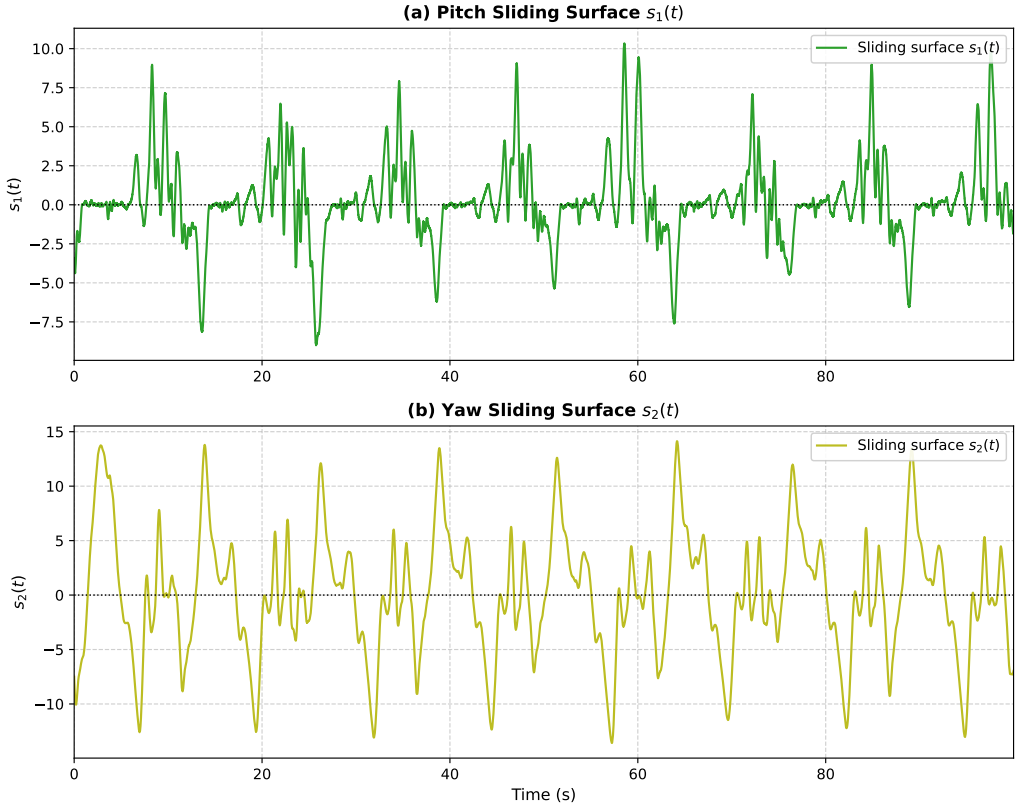


Figure 10. Real-time evolution of the B-STA sliding surfaces. (**a**) Pitch surface s (rad/s^2). (**b**) Yaw surface s_{yaw} (rad/s^2). Both surfaces remain bounded throughout, consistent with the practical stability result of Theorem 1. The residual amplitude reflects the composite construction (24), in which the dominant term e_3 carries units of angular acceleration and is therefore numerically large relative to the tracking error itself, and the larger residual on yaw is consistent with the bound evaluated for that axis in Table 2.

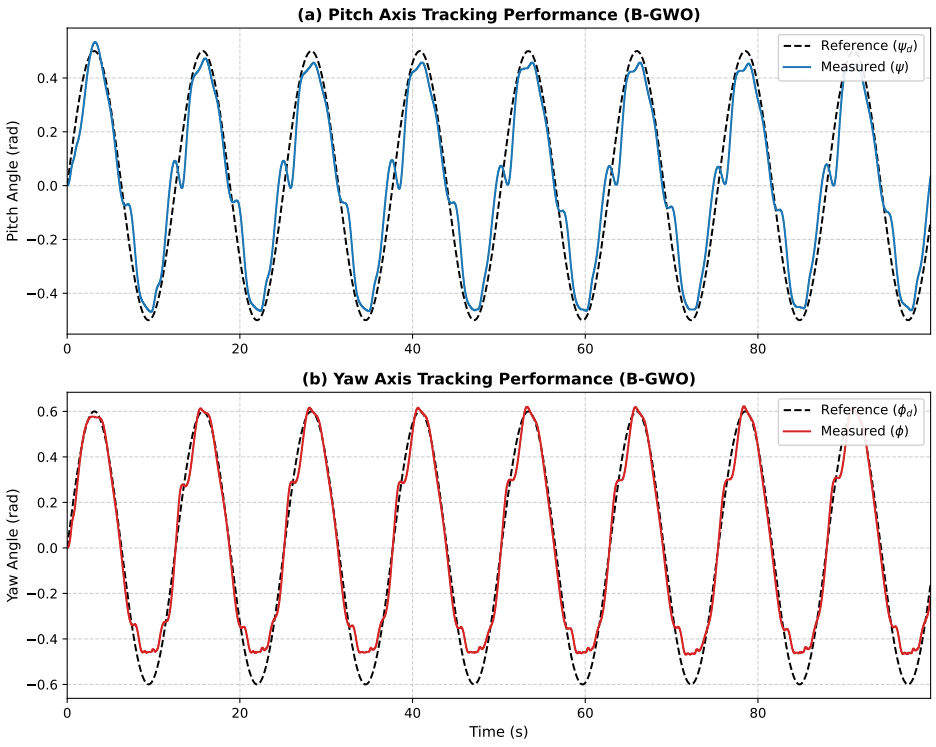


Figure 11. Real-time B-GWO tracking. (**a**) Pitch angle ψ (rad). (**b**) Yaw angle ϕ (rad).

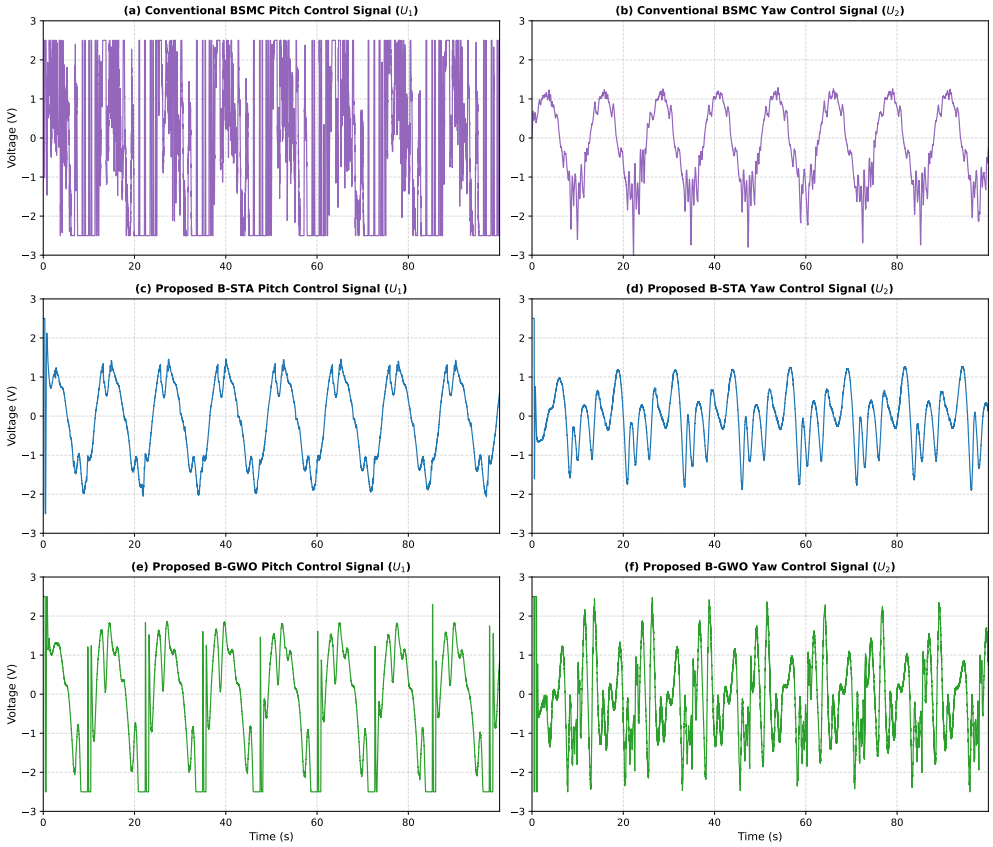


Figure 12. Applied control voltages over the 100 s window. (**a**) BSMC U_1 (V). (**b**) BSMC U_2 (V). (**c**) B-STA U_1 (V). (**d**) B-STA U_2 (V). (**e**) B-GWO U_1 (V). (**f**) B-GWO U_2 (V).

8.5. Discussion

8.5.1. Tracking Accuracy

B-STA attains the lowest normalised tracking error on both axes, at 7.94% of the commanded amplitude on pitch and 10.43% on yaw. The comparison with B-GWO must be made on normalised rather than absolute error, because the two runs commanded different excursions: the B-GWO yaw amplitude was 0.60 rad against 1.00 rad for B-STA. On absolute RMSE the B-GWO yaw figure is the lower of the two, but this reflects the smaller commanded excursion and the ordering reverses once normalised. We note this explicitly because it would be easy, and misleading, to read an axis-dependent performance split from the absolute column alone.

On pitch the margin is substantial in both columns. The mechanism is the super-twisting integrator state, which supplies continuous compensation for the gravitational restoring torque $M_g \sin x_1$. This torque is the largest single term in the pitch model at moderate angles and is persistent rather than oscillatory, so a controller without integral action carries a standing offset that RMSE registers directly and that ITAE penalises through its time weighting. Section 7.3 isolates the same mechanism under a controlled matched disturbance.

Margins against each comparator.

Table 7 supports three separate comparisons on each axis, and the differences are of different magnitudes in each. On pitch, B-STA reduces root mean square error by 53.4% against the first-order BSMC baseline, by 60.2% against the online RS-QBHO scheme, and by 62.9% against B-GWO, with corresponding ITAE reductions of 60.5%, 68.1% and 69.5% ordered identically. The first two of these are absolute comparisons: as recorded in the caption of Table 7, the commanded amplitudes of the earlier campaign are not stated in [29]

and could not be confirmed from the archived logs, so those margins are reported for completeness and are not amplitude-corrected. The comparison with B-GWO is corrected, both runs being recorded in the present study.

That the ITAE margin exceeds the RMSE margin in all three comparisons is a consequence of the time weighting rather than a separate effect: a standing offset contributes to RMSE uniformly across the record but to ITAE with linearly increasing weight, so an index that penalises persistent error credits the integral action of the super-twisting term more heavily than an index that averages it.

On yaw the comparisons divide. Against the two controllers recorded in the earlier campaign the margins remain substantial, at 43.9% and 54.8% in RMSE and 53.0% and 60.5% in ITAE, subject to the same amplitude qualification. Against B-GWO the absolute figures reverse, B-GWO recording 0.0683 rad against 0.1043 rad, but the two runs commanded 0.60 rad and 1.00 rad respectively, and on the normalised figures the ordering is restored at 10.43% against 11.39%. This last comparison is the amplitude-corrected one, and it is the one on which the opening claim of this subsection rests.

The margin is axis-dependent, and the residual bounds predict its size.

The normalised improvement over B-GWO is 53.7% on pitch and 8.4% on yaw, a ratio of approximately six. Section 2.5 evaluates the residual perturbation entering each sliding variable from the coefficients of Table 1 and the saturation limit alone, and finds the yaw figure to exceed the pitch figure by a factor of approximately six in Δ_1, driven by the reaction-torque coefficient $1.75|K_c|/I_2$ and by the elevated yaw transfer gain. The two ratios agree, and the agreement is not fitted: the bounds of Table 2 were computed without reference to the recorded runs.

The mechanism is the one identified in Section 5.2. The super-twisting injection acts on the residual entering the sliding variable, so its advantage over the same recursive law without that injection is largest where the residual is smallest and the term therefore operates closest to the regime the gain condition admits. On pitch, where the dominant load is the gravitational restoring torque and the evaluated residual is two orders of magnitude smaller in Δ_0, the injection recovers most of the available margin. On yaw, where the reaction torque acts on one third the inertia and is amplified by the larger backstepping gains, it recovers considerably less. The trade-off recorded in Section 2.5, that the gains delivering the required yaw error convergence simultaneously amplify the residual reaching the yaw sliding variable, is therefore visible in the measured tracking margins and not only in the sliding surfaces of Figure 10.

Tracking accuracy is not purchased with actuator effort.

One figure from the remaining columns belongs here because it bears directly on the tracking result. On pitch, B-STA attains its 53.4% reduction in RMSE against the first-order baseline while simultaneously reducing control energy by 72.9% and total variation by 96.6%. The improvement in tracking is therefore not obtained at the cost of increased actuator activity, which is the trade a higher-gain design would ordinarily make. Section 8.5.3 treats the energy and actuator-load columns in full.

8.5.2. Interpretation of the ITAE Column

With the corrected baseline values, the ITAE and RMSE rankings agree throughout the table and no separate reconciliation between the two indices is required. B-STA reduces pitch ITAE by 60.5% relative to the first-order baseline, consistent with the RMSE reduction on the same axis and attributable to removal of the gravity-related standing offset described above.

We draw attention to the consistency check (53), since the discrepancy it exposes is easy to introduce and difficult to notice. Time-weighted integral indices scale with the assumed

sampling interval whereas root mean square error does not, so a mismatch between the rate at which a signal is logged and the interval used in post-processing corrupts one index while leaving the other intact. Any reported pair violating (53) cannot have been computed consistently. Every entry in Table 7 satisfies it, at between 0.59 and 0.81 of the ceiling.

On the magnitude of the reported values.

The absolute size of the ITAE figures warrants comment, because it invites a misreading. The index weights the absolute error by elapsed time, so for a fixed error level its value scales with $T^2/2$ rather than with T. Over the 100 s window used here, $\int_0^T t\,\mathrm{d}t = 5000\,\mathrm{s}^2$, so a mean absolute tracking error of only 0.02 rad produces an ITAE of approximately 100 rad·s on its own. The entries of Table 7 are of precisely this order and are therefore consistent with the tracking errors recorded alongside them. The apparent acceleration of the accumulation in the second half of each run reflects the linearly increasing weight and not a growth in the underlying error, which Figures 9 and 11 show to be stationary over the window.

Two further properties follow and are stated to forestall misinterpretation. The index is not normalised by the horizon, so values are comparable within Table 7, where every run has the same duration, but not against figures obtained over a different window. And the reported values do not accumulate error across the stages of the design. Each entry is formed from the measured tracking error of a single axis at the plant output, e_1 on pitch and e_4 on yaw. The three recursive steps of Section 4 and the three-stage cascade of Section 5 are a decomposition used for synthesis and for analysis; they are not a signal path along which residuals compound into the reported index, and the two axes are scored independently of one another.

8.5.3. Control Energy and Actuator Load

B-STA records the lowest control energy on both axes. On pitch, energy falls by 72.9% relative to the first-order baseline, from 390.47 to 105.94 V^2·s. On yaw it falls by 49.5% relative to the baseline and by 56.2% relative to B-GWO. These are the largest and most consistent margins in the table.

The total variation column records a pronounced difference on the pitch axis, where the baseline accumulates 2416.7 V of cumulative excursion against 82.4 V for B-STA, and spends 48.9% of the run at the saturation limit against 0.27%. We report this as a measured property of the recorded runs and are deliberately cautious about its attribution. Simulation of the same two control laws on the nominal coupled model does not reproduce the baseline behaviour: the simulated baseline reaches the saturation limit for 0.28% of the run. The hardware behaviour therefore involves effects absent from the design model, most plausibly stiction, cable tension and amplifier dynamics on the gravity-loaded axis, interacting with the equivalent-control term. The reduction is real and is measured, but it should not be attributed to the injection structure alone.

Reduced high-frequency content in the actuator command is expected to lower brush wear, winding dissipation and vibratory loading of the transmission. Quantification of that effect requires accelerated life testing and lies outside the scope of the present experiments.

8.5.4. Behaviour of the Online Scheduler

B-GWO tracks both axes without discontinuous injection and without instability, confirming that a population-based scheduler can be operated inside a real-time loop on this platform. Its normalised tracking accuracy does not match B-STA on either axis, and its yaw control energy is the highest in the table, which is consistent with the gain excursions produced by successive test windows.

Section 6.5 identifies a specific reason those excursions persist rather than decaying. At $T_s = 0.01$ s a generation occupies 10 s, so the 100 s run completes ten generations against the $G_{\max} = 20$ for which the exploration coefficient was scheduled. The search therefore operates in its exploratory regime throughout and never contracts toward the incumbent leaders. The measured B-GWO performance should accordingly be read as the performance of an exploration phase rather than of a converged search, and the natural correction is to match $G_{\max}$ to the generations the horizon admits, or to lengthen the horizon. We record this as a configuration lesson of the same practical character as Proposition 1: an online scheduler is constrained not only by how its fitness is formed but by whether the experiment is long enough for its schedule to run to completion.

The result of Proposition 1 bears directly on how such schemes must be configured. In a preliminary implementation the candidate fitness was evaluated at the current sample rather than accumulated over the test window. Under that configuration the applied gains remained at the regularisation centre $\bar{\boldsymbol{\theta}}$ for the entire run, varying by less than 2% of the gain box width, exactly as Proposition 1 predicts. The windowed formulation (49) is what makes the search respond to the plant, and this is the substantive content of the test-and-settle mechanism rather than the scheduling convenience it might appear to be.

The choice of gain box is equally consequential. In a configuration confining the yaw scheduler to the pitch box, which caps k_4 at five against the value of 20 that the fixed-gain design requires on that axis, yaw tracking degraded by roughly an order of magnitude. Widening the yaw box as in Table 4 restored performance. An online scheduler cannot recover from a search space that excludes the operating point, and the box must therefore be set from the same physical reasoning that governs fixed-gain tuning.

The pitch box furnishes a second instance of the same condition, and we record it because it bears on how the pitch row of Table 7 should be read. Table 4 caps k_1 at five, whereas the fixed-gain design of Table 3 uses $k_1 = 8$ and $k_2 = 6$ on that axis, so the pitch box also excludes the operating point to which its own channel was tuned. Figure 6 shows k_1 held at the box ceiling over extended intervals, which is the signature of a search pressing against a bound rather than settling within one, and B-GWO records its worst normalised error on precisely that axis, 17.16% against 11.39% on yaw. The condition is therefore evidenced by both channels rather than by the yaw configuration alone, and the pitch figures reported for B-GWO should be read as those of a scheduler operating under a box that excludes the loop gain the channel requires.

8.5.5. On Axis-Dependent Controller Assignment

The submitted version of this work proposed a mixed assignment placing B-STA on pitch and B-GWO on yaw, inferred from the absolute RMSE column. The normalised comparison in Table 7 does not support that inference: once the differing reference amplitudes are accounted for, B-STA is the better performer on both axes. We therefore make no axis-dependent design recommendation. A mixed architecture was in any case never implemented as a separate closed loop and would require dedicated validation before any such claim could be advanced.

9. Conclusions

This paper has presented an observer-based output-feedback architecture for the Twin Rotor MIMO System and validated it on physical hardware against baselines recorded on the same real-time interface.

Three results are established. First, the coupled six-state TRMS model is not in strict-feedback form, so recursive backstepping cannot be applied to it directly. The decentralised decomposition adopted here is a structural requirement rather than a simplification, and

the coupling it discards is shown to be unmatched with respect to the original state chain but matched with respect to the composite sliding variable, which is what places it in the channel the super-twisting term can reject. The observer, subject to no such structural restriction, retains the coupling in full.

Second, the closed loop is analysed as a three-stage cascade covering observer error, sliding variable and tracking error. The reduced tracking dynamics are input-to-state stable with respect to the sliding variable with an explicit gain, and the relation on which this rests holds throughout the reaching phase rather than only on the sliding set. The result is practical stability with an ultimate bound expressed in terms of the residual perturbation, which is derived from the coupling coefficients and the actuator saturation limit rather than assumed. Evaluating that bound numerically shows the perturbation entering the yaw sliding variable to exceed the pitch figure by approximately a factor of six, and to exceed the range admitted by the super-twisting gain condition on both axes. We report this rather than adjust the gains to fit the condition, and we show that the resulting prediction, namely that the yaw surface will not reach the tight residual neighbourhood the pitch surface reaches, is borne out in both simulation and experiment.

Third, on a 100 s experimental run the B-STA controller attained the lowest normalised tracking error on both axes, at 7.94% and 10.43% of the commanded amplitude, together with the lowest control energy on both axes, reducing pitch control energy by 72.9% relative to the first-order baseline. Numerical disturbance rejection tests on the fully coupled model confirm the underlying mechanism directly: under a matched actuator step the super-twisting integrator state migrates to a new steady level that cancels the disturbance, leaving a post-event pitch error of 0.0002 rad, whereas the same recursive law without the second-order injection retains a permanent offset of 0.282 rad and does not recover.

A separate methodological result concerns online metaheuristic gain scheduling. Proposition 1 establishes that a fitness evaluated at a single sample is common to all candidates and cancels from the population ranking, so the search reduces to the minimiser of its own regularisation term and produces a controller that is nominally adaptive and effectively fixed. Windowed evaluation, in which each candidate is applied to the plant and scored on its own accumulated error, is therefore a requirement for well-posedness rather than a scheduling convenience. Our experiments add two practical conditions. The gain search space must contain the operating point that the physical channel requires, since a scheduler cannot recover from a box that excludes it, a condition on which both axes of the present configuration bear. And the experimental horizon must be long enough to complete the generation schedule, since a run that exhausts only half of it reports the behaviour of an exploration phase rather than of a converged search.

Several limitations are stated plainly. All experimental results correspond to a single run per controller with no multi-trial statistics, which is a particular limitation for a stochastic scheduler. The disturbance rejection study is numerical rather than experimental, because repeatable disturbance injection on this platform would require instrumentation that was not available. The safety provisions on the scheduler bound candidate excursions but do not constitute a stability certificate for the time-varying-gain plant, and no such certificate is claimed. The super-twisting gains of Table 3 do not satisfy the sufficient condition of Lemma 1 at the evaluated residual bounds, so the guarantee available on the reported runs is boundedness rather than finite-time arrival at a tight residual set. The B-GWO gain boxes of Table 4 do not contain the fixed-gain operating point of either channel, so the scheduler results are those of a search constrained below the loop gain the plant requires. On the yaw axis under nominal simulated conditions the first-order baseline outperforms B-STA, and we report this rather than confine the comparison to the axis that favours the proposed method.

Future work will address multi-trial statistical characterisation of the scheduler, experimental disturbance injection, frequency-domain quantification of actuator load to complement the total-variation measure reported here, and tighter residual estimation through disturbance observers operating alongside the state observer. The yaw channel in particular warrants a dedicated study of the trade-off identified in Section 5.2, in which the backstepping gains that deliver the required error convergence simultaneously amplify the residual coupling reaching the sliding variable, and for which a disturbance observer compensating the reaction torque directly is the most promising route.

Author Contributions: Conceptualization, A.B. and A.K.; methodology, A.B., A.K. and S.T.; software, A.B. and A.K.; validation, A.B., A.K., S.T., A.S. and A.A.; formal analysis, A.K., S.T. and M.H.Z.; investigation, A.B. and M.H.Z.; resources, A.S., A.A. and J.-J.T.; data curation, A.B. and A.A.; writing, original draft preparation, A.B. and A.K.; writing, review and editing, A.B., A.K., S.T., A.S., A.A., M.H.Z. and J.-J.T.; visualisation, A.B., A.K. and S.T.; supervision, A.K. and A.S.; project administration, A.K. and J.-J.T.; funding acquisition, J.-J.T. All authors have read and agreed to the published version of the manuscript.

Funding: The Article Processing Charge was funded by Multimedia University, Malaysia.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The experimental data, the controller and observer source files, the disturbance study script, and the metric computation scripts supporting the reported results are available from the corresponding author on request.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

TRMS	Twin Rotor MIMO System
B-STA	Backstepping–Super-Twisting Algorithm
B-GWO	Backstepping–Grey Wolf Optimiser
BSMC	Backstepping–Sliding Mode Control
STA	Super-Twisting Algorithm
GWO	Grey Wolf Optimiser
RS-QBHO	Rate-Constrained Sequential Quantum Black Hole Optimisation
RMSE	Root Mean Square Error
ITAE	Integral of Time-weighted Absolute Error
CE	Control Effort
TV	Total Variation
SMC	Sliding Mode Control
ISS	Input-to-State Stable
VTOL	Vertical Take-Off and Landing

References

1. Rahideh, A.; Shaheed, M.H. Mathematical dynamic modelling of a twin-rotor multiple input-multiple output system. *Proc. Inst. Mech. Eng. Part I J. Syst. Control Eng.* **2007**, *221*, 89–101. [CrossRef]
2. Pratap, B.; Purwar, S. Real time optimal control of a twin rotor MIMO system with state dependent Riccati equation. *Int. J. Autom. Comput.* **2012**, *9*, 195–202.
3. Raghavan, R.; Thomas, S. Practically implementable model predictive controller for a twin rotor multi-input multi-output system. *J. Control Autom. Electr. Syst.* **2017**, *28*, 170–180. [CrossRef]
4. Khalil, H.K. *Nonlinear Systems*, 3rd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 2002.

5. Krstic, M.; Kanellakopoulos, I.; Kokotovic, P.V. *Nonlinear and Adaptive Control Design*; John Wiley & Sons: New York, NY, USA, 1995.
6. Levant, A. Sliding order and sliding accuracy in sliding mode control. *Int. J. Control* **1993**, *58*, 1247–1263. [CrossRef]
7. Moreno, J.A.; Osorio, M. Strict Lyapunov functions for the super-twisting algorithm. *IEEE Trans. Autom. Control* **2012**, *57*, 1035–1040. [CrossRef]
8. Mondal, S.; Mahanta, C. Adaptive second-order sliding mode controller for a twin rotor multi-input-multi-output system. *IET Control Theory Appl.* **2012**, *6*, 2157–2167. [CrossRef]
9. Van, M.; Ge, S.S. Adaptive fuzzy integral sliding-mode control for robust fault-tolerant control of robot manipulators with disturbance observer. *IEEE Trans. Fuzzy Syst.* **2021**, *29*, 1284–1296. [CrossRef]
10. Beltran, B.; Benbouzid, M.E.H.; Ahmed-Ali, T. Second-order sliding mode control of a doubly fed induction generator driven wind turbine. *IEEE Trans. Energy Convers.* **2012**, *27*, 261–269. [CrossRef]
11. Chen, G.; Wu, W.; Yang, C.; Hu, H.; Zhang, J.; Shi, J.; Wu, C.; Ma, Y. Practical finite-time observer-based adaptive backstepping super-twisting sliding mode control for a deep-sea hydraulic manipulator. *IEEE Trans. Ind. Electron.* **2025**, *72*, 9569–9579. [CrossRef]
12. Guan, C.; Pan, S. Adaptive backstepping sliding mode control for direct driven hydraulics. *Energies* **2020**, *13*, 5635.
13. Wang, W.; Tan, F.; Wu, J.; Ge, H.; Wei, H.; Zhang, Y. Adaptive integral backstepping controller for PMSM with AWPSO parameters optimization. *Energies* **2019**, *12*, 2596. [CrossRef]
14. Gao, B.; Liu, Y.J.; Liu, L. Adaptive neural fault-tolerant control of a quadrotor UAV via fast terminal sliding mode. *Aerosp. Sci. Technol.* **2022**, *129*, 107818. [CrossRef]
15. Singh, P.; Giri, D.K.; Ghosh, A. Robust sliding mode augmenting integral backstepping control of a highly maneuverable aircraft. *Aerosp. Sci. Technol.* **2021**, *119*, 107223.
16. Beltran, B.; Ahmed-Ali, T.; Benbouzid, M.E.H. Sliding mode power control of variable-speed wind energy conversion systems. *IEEE Trans. Energy Convers.* **2008**, *23*, 551–558. [CrossRef]
17. Cucuzzella, M.; Incremona, G.P.; Ferrara, A. Design of robust self-configuring controllers for AC microgrids. *IEEE Trans. Smart Grid* **2019**, *10*, 1509–1520.
18. Frikha, A.; Benômar, Y.; El Baghdadi, M.; Hegazy, O. Multiphase motors and drive systems for electric vehicle powertrains: State of the art analysis and future trends. *Energies* **2023**, *16*, 768. [CrossRef]
19. Lin, F.J.; Chen, S.Y. Robust control of linear synchronous motor using adaptive backstepping sliding-mode control. *IET Electr. Power Appl.* **2015**, *9*, 556–565.
20. Juang, J.G.; Lin, R.W.; Liu, W.K. Intelligent control for a twin-rotor multi-input multi-output system. *Eng. Appl. Artif. Intell.* **2008**, *21*, 1024–1034.
21. Rahideh, A.; Shaheed, M.H. Constrained nonlinear predictive control for a twin rotor MIMO system. *Proc. Inst. Mech. Eng. Part G J. Aerosp. Eng.* **2012**, *226*, 175–187.
22. Zaihidee, F.M.; Mekhilef, S.; Mubin, M. Robust control of twin rotor MIMO system using fractional order sliding mode control. *Ain Shams Eng. J.* **2019**, *10*, 411–421.
23. Ibraheem, I.K.; Ibraheem, G.A. Fractional order PID control design for a twin rotor MIMO system. *J. Eng. Sci. Technol. Rev.* **2020**, *13*, 98–108.
24. Ramya, A.; Imthiaz, S.A. Design of PSO based PID controller for Twin Rotor MIMO System. In Proceedings of the International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT), Chennai, India, 3–5 March 2016; pp. 433–438.
25. Guan, C.; Pan, S. Nonlinear adaptive position tracking of an electro-hydraulic actuator based on backstepping. *Mech. Syst. Signal Process.* **2020**, *134*, 106293.
26. Mirjalili, S.; Mirjalili, S.M.; Lewis, A. Grey wolf optimizer. *Adv. Eng. Softw.* **2014**, *69*, 46–61. [CrossRef]
27. Precup, R.E.; David, R.C.; Petriu, E.M. Grey wolf optimizer-based approaches to PID controller tuning. *Appl. Soft Comput.* **2017**, *55*, 397–407.
28. Mirjalili, S.; Saremi, S.; Mirjalili, S.M.; Coelho, L.S. Multi-objective grey wolf optimizer: A novel algorithm for multi-criterion optimization. *Expert Syst. Appl.* **2016**, *47*, 106–119. [CrossRef]
29. Kacimi, A.; Beloufa, A.; Tahraoui, S.; Senoussaoui, A.; Zaid, M.H.; Azzouz, A.; Tiang, J.-J. Robust Control of Twin Rotor MIMO Systems Under Unmodeled Dynamics: Comparative Experimental Validation of Hybrid BSMC and Online QBHO Strategies. *Actuators* **2026**, *15*, 236. [CrossRef]
30. Kacimi, A.; Mostefaoui, M.; Beloufa, A.; Tahraoui, S.; Azzouz, A.; Tiang, J.-J.; Zaid, M.H. High-Gain Observer-Based Backstepping Control for Real-Time Trajectory Tracking of a Twin Rotor MIMO System: Adaptive Tuning Functions Versus Metaheuristic Gain Optimization. *Actuators* **2026**, *15*, 411. [CrossRef]
31. Feedback Instruments Ltd. *Twin Rotor MIMO System Control Experiments Manual 33-949S*; Feedback Instruments Ltd.: Crowborough, UK, 2006.

32. Beloufa, A.; Tahraoui, S.; Kacimi, A.; Allouache, H.; Tiang, J.-J.; Azzouz, A.; Zaid, M.H. Robust Backstepping Control of a Twin Rotor MIMO System via an RBF-Tuned High-Gain Observer. *Automation* **2026**, *7*, 40. [CrossRef]
33. Bornard, G.; Hammouri, H. A high gain observer for a class of uniformly observable systems. In Proceedings of the 30th IEEE Conference on Decision and Control, Brighton, UK, 11–13 December 1991; pp. 1494–1496.
34. Beloufa, A.; Tahraoui, S.; Kacimi, A.; Allouach, H.; Tiang, J.-J.; Azzouz, A. Experimental Validation of Robust Backstepping Control for TRMS Using an Interval Type-2 Fuzzy Observer. *Eng* **2026**, *7*, 171. [CrossRef]
35. Gokgan, Y.C.; Efe, M.O. Flight control of a quadrotor using super twisting sliding mode control. *Aerosp. Sci. Technol.* **2020**, *98*, 105658. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.