actuators

MDPI

Article

Chattering-AwareFuzzy Gain Scheduling of a High-Gain Observer for a Twin Rotor MIMO System: Interval Type-3 Formulation, Neural Benchmarks and Real-Time Validation

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

[1] LIS2T Laboratory, Department of Electrotechnics, University of Mustapha Stambouli, Mascara 29000, Algeria; s.abderrahmane@univ-mascara.dz
[2] Laboratory of Electrical Engineering and Renewable Energies (LGEER), Electronic Department, Hassiba Benbouali University, Chlef 02000, Algeria; a.beloufa@univ-chlef.dz
[3] Department of Instrumentation Maintenance, Institute of Maintenance and Industrial Safety, University of Oran 2, Oran 31000, Algeria; kacimi.abderrahmane@univ-oran2.dz
[4] Laboratory of Electronics, Advanced Signal Processing, and Microwave (LESM), Department of Telecommunications, Saida University, Saida 20000, Algeria; abdelbasset.azzouz@univ-saida.dz
[5] 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
[6] 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

Output-feedback control of the Twin Rotor MIMO System (TRMS) relies on a high-gain observer (HGO) whose gain trades estimation speed against the amplification of encoder quantisation into the control voltage. This paper proposes a chattering-aware fuzzy scheduler that adapts the HGO gain online from three measured features: the filtered estimation error above an online noise-floor estimate, its trend, and a chattering index of the control voltage. The scheduler is formulated with interval type-3 (IT3) sets evaluated through three alpha-planes, embedded in a corrected HGO with saturated output injection and a decentralised backstepping controller, and analysed with a practical-stability condition for the time-varying gain that is evaluated along the recorded trajectories. A multilayer perceptron trained with ADAM and a radial basis function network trained with normalised least mean squares, which share the same features and gain mapping, serve as neural benchmarks. In real-time experiments at 1 kHz with simultaneous pitch and yaw tracking and a common 10 ms scheduler period, the fuzzy scheduler reduced the pitch tracking RMSE by 35% and the RMS estimation error by 22 to 32% relative to both neural schedulers ($p < 0.001$ over six reference cycles), at the cost of a 10 to 18% higher control total variation. Ablations on offline replays of the recorded data and quantised Monte Carlo simulations show that removing the chattering input increases the total variation by up to 100%, that a fixed gain equal to the mean scheduled gain is worse in both estimation accuracy and smoothness under model mismatch, and that reducing the IT3 sets to interval type-2 sets changes the results by at most 0.2%. The benefit therefore stems from the chattering-aware scheduling information rather than from the third dimension of the fuzzy sets. In a compiled implementation, 99% of the scheduler updates take less than 3 µs.

Keywords: twin rotor MIMO system; high-gain observer; gain scheduling; chattering; interval type-3 fuzzy logic; backstepping control; neural networks; real-time experiments

Check for updates

Academic Editor: Guanghong Yang

Received: 5 November 2026
Revised: 16 November 2026
Accepted: 20 November 2026
Published: 23 November 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.

1. Introduction

The Twin Rotor MIMO System (TRMS) is a laboratory platform that reproduces the essential flight mechanics of a helicopter with two degrees of freedom: a pitch motion driven by a main rotor and a yaw motion driven by a tail rotor. Its dynamics combine gravitational loading on the vertical axis, friction-dominated motion on the horizontal axis, first-order rotor dynamics, gyroscopic coupling and a strong aerodynamic cross-coupling between the two rotors [1,2]. These features, together with a systematic mismatch between the phenomenological model and the physical rig, make the TRMS a demanding benchmark for nonlinear control and estimation, and it continues to be used to validate model-based and learning-based flight controllers in real time [3–7].

Backstepping provides a systematic recursive design with Lyapunov guarantees for systems in strict-feedback form [8]. On the TRMS, however, only the pitch and yaw angles are measured by optical encoders, whereas a backstepping law requires the angular rates and the rotor torques. Output feedback is therefore usually obtained with a high-gain observer (HGO) [9–11]. For uniformly observable nonlinear systems, the HGO of Gauthier, Hammouri and Othman [12] and its multi-output extension [13] provide exponential convergence at a rate fixed by a single scalar gain θ, and the separation principle guarantees that the output-feedback loop recovers the state-feedback performance when θ is large enough [11]. The price is a well-known sensitivity to measurement noise: the i-th estimated derivative amplifies the measurement perturbation by a factor of order θ^{i-1}, so that the gain that minimises the estimation error is bounded from above by the noise level [14]. On the TRMS the dominant perturbation is the encoder quantisation, which the third observer state, the rotor torque, receives with an amplification of order θ^2 and transfers directly into the control voltage.

Several remedies have been proposed. Switched-gain observers use a high gain during transients and a low gain in steady state [15]; gain-adaptation laws drive θ from the output error [16,17]; nonlinear injection and limited-gain-power structures reduce the sensitivity of the high-order terms [18,19]. Our previous works on the same rig replaced the constant gain by an interval type-2 (IT2) fuzzy supervisor [20] and by an online radial basis function (RBF) tuner [21], and showed that adapting the observer gain is decisive for real-time stability. Those schedulers, however, observed only the output-estimation or tracking error. They did not measure the quantity that a high gain degrades, namely the chattering of the control voltage, and they did not distinguish an estimation error caused by a transient from one caused by sensor resolution.

Type-3 fuzzy logic offers a general language for this problem. In an interval type-3 (IT3) fuzzy set, the secondary membership is itself an interval type-2 set, which adds a third dimension of uncertainty to the footprint of uncertainty of type-2 sets [22,23]. IT3 fuzzy systems have shown better robustness than their type-1 and type-2 counterparts in identification, estimation and control under noisy conditions [22,24–26]. Their computation relies on alpha-plane decompositions [27,28] combined with Karnik–Mendel-type reduction [29,30]. Whether the third dimension brings a measurable benefit in a given application is, however, an empirical question, and it is examined explicitly in this work.

The scheduling problem considered here is an instance of a broader trade-off in actuator-based control between tracking accuracy, actuation effort, energy and signal smoothness. In vibration suppression, passive and active piezoelectric solutions for satellite panels have been compared in terms of damping performance and actuation demand [31], and integrated active and passive isolators combine the reliability of passive absorption with the controllability of active actuation [32]. In electric drives and power converters, sliding-mode and observer-based designs address the same trade-off: barrier-function adaptation limits the switching gains of super-twisting control and hence its chattering [33],

super-twisting extended state observers provide disturbance estimates without high-gain amplification of noise [34], and fixed-time second-order sliding modes provide robust voltage regulation under parameter variations [35]. In the present problem, the accuracy and smoothness objectives meet in a single scalar, the observer gain, and the scheduler must arbitrate between them online from measured signals.

Table 1 positions the present work with respect to our closest previous studies on the same rig. Reference [6] keeps the observer gain constant and adapts the backstepping gains, whereas [20,21] schedule the observer gain from the error signal only, at a loop rate of 100 Hz, and rely on the separation principle established for a constant gain.

Table 1. Comparison of the present work with previous studies on the same TRMS rig.

	Ref. [6]	**Ref. [20]**	**Ref. [21]**	**This Work**
Observer	HGO with constant gains ($\theta_\psi = 2$, $\theta_\phi = 1.5$)	HGO, gain $\theta = 20\,\Psi$	HGO, $\theta_\psi \in [1, 5]$, $\theta_\phi \in [0.5, 2]$, low-pass filtered	Corrected HGO, consistent scaling, saturated injection, $\theta_\psi \in [4, 25]$, $\theta_\phi \in [3, 20]$
Adapted quantity	Backstepping gains (tuning functions, BSO, ABC)	Observer gain	Observer gain	Observer gain
Adaptation mechanism	Adaptive laws and online metaheuristics	IT2 Takagi–Sugeno fuzzy supervisor, 25 rules	RBF network, 7 units, gradient descent towards a fixed target	Chattering-aware fuzzy scheduler, 18 rules; MLP-ADAM and RBF-NLMS benchmarks on a performance cost
Scheduling inputs	Not applicable	Estimation error and its rate	Normalised tracking error	Excess estimation error above an online noise floor, its trend, chattering index of U
Loop and update rate	1 kHz	100 Hz	100 Hz	1 kHz loop, scheduler every 10 ms
Theoretical result	Lyapunov analysis with a constant observer gain	Separation principle, constant gain	Separation principle, constant gain	Practical stability with a time-varying, rate-limited gain, evaluated along recorded data
Experimental evaluation	Four architectures, references of 0.5 and 1.0 rad	Fixed versus fuzzy gain; RMSE 0.0883 and 0.1979 rad	Fixed versus RBF gain; RMSE 0.0385 and 0.1080 rad	Three schedulers, three ablations, cycle-wise statistics, offline replays and Monte Carlo simulations; RMSE 0.0658 and 0.0599 rad

RMSE values are those reported for pitch and yaw in each study; reference signals and loop rates differ between studies, so these values indicate orders of magnitude rather than a controlled comparison.

This paper addresses these gaps. Its contributions are as follows:

1. A chattering-aware fuzzy scheduler for the HGO gain, denoted IT3-CA, is proposed. It uses three features that are measured online: the filtered output-estimation error above an adaptive noise-floor estimate, the trend of this excess error, and a chattering index of the control voltage. Its 18-rule base raises the gain when the estimation error grows above the sensor resolution and lowers it when the control signal starts to chatter. The scheduler is formulated with IT3 sets, and an ablation shows that its benefit stems from the scheduling information rather than from the third dimension of the sets.

2. The scheduler is embedded in a corrected HGO built on the coupled six-state model, with a consistent gain scaling and a saturated output injection that bounds peaking. A practical-stability condition is derived for the time-varying, rate-limited gain, with an optimised Lyapunov matrix, an averaged-rate formulation and an evaluation of the condition along the recorded operating points.
3. Two neural schedulers are designed as a fair benchmark: a multilayer perceptron trained with ADAM (MLP-ADAM) and an RBF network trained with normalised least mean squares (RBF-NLMS). Both share the features, gain mapping and rate-limiter structure of the fuzzy scheduler and learn online from a performance cost rather than from a fixed target.
4. The three schedulers are compared on the physical TRMS at 1 kHz, in two-degree-of-freedom operation with simultaneous pitch and yaw tracking and a common scheduler period of 10 ms, using cycle-wise statistics. The comparison is complemented by offline replays of the recorded data, which isolate the contributions of the chattering input, the third fuzzy dimension and the gain adaptation itself, by quantised Monte Carlo simulations under nominal and mismatched conditions, and by measured execution times.

The remainder of the paper is organised as follows. Section 2 presents the model, the controller, the observer, the stability condition, the three schedulers and the experimental setup. Section 3 reports the simulation and experimental results. Section 4 discusses the trade-offs and the limitations of the study, and Section 5 concludes the paper.

2. Materials and Methods

2.1. TRMS Model

The TRMS consists of a beam pivoted on a tower, with a main rotor at one end and a tail rotor at the other, and a counterweight that sets the static equilibrium of the pitch axis [2]. As illustrated in Figure 1, the main rotor turns in a plane parallel to the beam axis and produces a thrust normal to the beam, which rotates the beam about the horizontal axis (pitch angle ψ) against the gravitational moment of the beam and counterweight assembly. The tail rotor turns in a vertical plane that contains the beam and produces a horizontal thrust, which rotates the whole assembly about the vertical axis of the tower (yaw angle ϕ). Each rotor is driven by a DC motor whose voltage, U_1 for the main rotor and U_2 for the tail rotor, is the control input. The two angles are measured by incremental optical encoders located at the pivot and at the tower. The axes are coupled in both directions: the reaction torque of the main rotor acts on the yaw axis, and the yaw rate produces a gyroscopic moment on the pitch axis.

With the state vector $\mathbf{x} = [x_1, \ldots, x_6]^T = [\psi,\ \dot{\psi},\ \phi,\ \dot{\phi},\ \tau_1,\ \tau_2]^T$, where ψ is the pitch angle, ϕ the yaw angle, and τ_1, τ_2 the main and tail rotor torques, the phenomenological model used for design reads [2,20]

$$\begin{cases} \dot{x}_1 = x_2, \\ \dot{x}_2 = \dfrac{1}{I_1}\Big(F_1(x_5) - M_g \sin x_1 - B_{1\psi} x_2 - K_{gy}\, x_4 \cos x_1\, F_1(x_5)\Big), \\ \dot{x}_3 = x_4, \\ \dot{x}_4 = \dfrac{1}{I_2}\Big(F_2(x_6) - B_{1\phi} x_4 - 1.75\, K_c\, F_1(x_5)\Big), \\ \dot{x}_5 = \dfrac{1}{T_{11}}\big(-T_{10} x_5 + k_1 U_1\big), \\ \dot{x}_6 = \dfrac{1}{T_{21}}\big(-T_{20} x_6 + k_2 U_2\big), \end{cases} \tag{1}$$

with the rotor maps $F_1(x_5) = a_1 x_5^2 + b_1 x_5$ and $F_2(x_6) = a_2 x_6^2 + b_2 x_6$, the control voltages $U_1, U_2 \in [-2.5, 2.5]$ V and the measured output $y = [x_1,\ x_3]^T = [\psi,\ \phi]^T$. The Coulomb friction terms of the manufacturer model are not retained in (1) and act as bounded disturbances. The parameters are listed in Table 2.

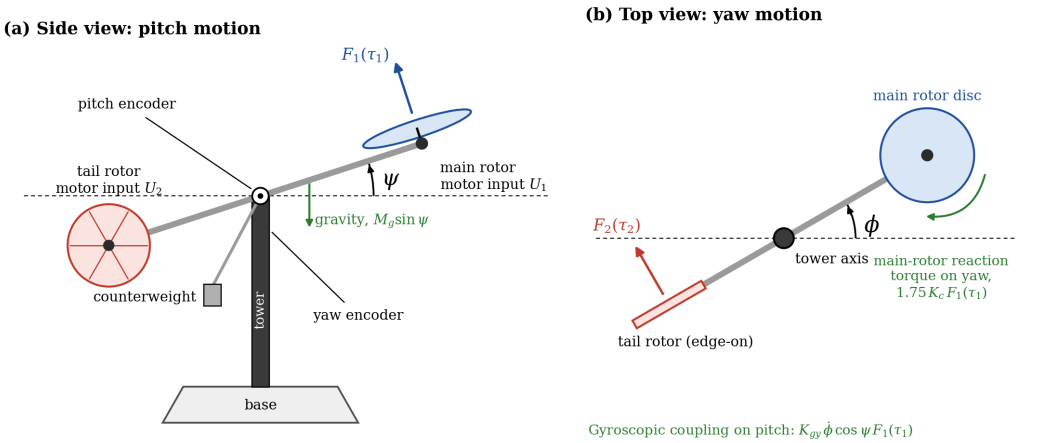

Figure 1. Schematic of the Twin Rotor MIMO System. (**a**) Side view: the main-rotor thrust $F_1(\tau_1)$ drives the pitch angle ψ against gravity. (**b**) Top view: the tail-rotor thrust $F_2(\tau_2)$ drives the yaw angle ϕ; the main-rotor reaction torque and the gyroscopic moment couple the two axes.

Table 2. TRMS model parameters [2].

Parameter	Value	Parameter	Value
I_1 (kg m^2)	6.8×10^{-2}	I_2 (kg m^2)	2×10^{-2}
a_1, b_1	0.0135, 0.0924	a_2, b_2	0.01, 0.09
M_g (N m)	0.32	K_{gy}, K_c	0.05, -0.2
$B_{1\psi}$, $B_{2\psi}$	6×10^{-3}, 1×10^{-3}	$B_{1\phi}$, $B_{2\phi}$	1×10^{-1}, 1×10^{-2}
k_1, T_{11}, T_{10}	1.1, 1.1, 1	k_2, T_{21}, T_{20}	0.8, 1, 1

2.2. Decentralised Backstepping Controller

Because (1) is not in strict-feedback form, the controller is designed on two decoupled subsystems, and the cross-coupling terms are treated as bounded disturbances [6,20]. The pitch design model is

$$\begin{aligned} \dot{x}_1 &= x_2, \qquad \dot{x}_5 = \frac{-T_{10}x_5 + k_1 U_1}{T_{11}}, \\ I_1 \dot{x}_2 &= G_1 := F_1(x_5) - M_g \sin x_1 - B_{1\psi} x_2 - B_{2\psi} x_1, \end{aligned} \tag{2}$$

where the linear term $B_{2\psi}x_1$ is kept as a design approximation of the residual restoring effects. Let $\psi_d(t)$ be a reference with bounded derivatives up to the third order. The tracking errors and the virtual controls are

$$e_1 = x_1 - \psi_d, \quad \alpha_1 = -k_{p1} e_1 + \dot{\psi}_d, \quad e_2 = x_2 - \alpha_1, \tag{3}$$

$$\begin{aligned} \alpha_2 &= -k_{p2} e_2 - e_1 + \frac{M_g \sin x_1 + B_{2\psi} x_1}{I_1} - k_{p1} x_2 + k_{p1} \dot{\psi}_d + \ddot{\psi}_d, \\ e_3 &= \frac{F_1(x_5) - B_{1\psi} x_2}{I_1} - \alpha_2. \end{aligned} \tag{4}$$

Differentiating α_2 along (2) gives

$$\begin{aligned}\dot{\alpha}_2 = &-(k_{p1}+k_{p2})\frac{G_1}{I_1} + x_2\Big(-k_{p1}k_{p2} - 1 + \frac{M_g\cos x_1 + B_{2\psi}}{I_1}\Big)\\ &+ (k_{p1}k_{p2}+1)\dot{\psi}_d + (k_{p1}+k_{p2})\ddot{\psi}_d + \dddot{\psi}_d,\end{aligned} \tag{5}$$

and enforcing $\dot{e}_3 = -k_{p3}e_3 - e_2$ yields the pitch control law

$$U_1 = \mathrm{sat}_{2.5}\left\{\frac{I_1 T_{11}}{k_1 D_1}\Big(-k_{p3}e_3 - e_2 + \frac{B_{1\psi}G_1}{I_1^2} + \dot{\alpha}_2\Big) + \frac{T_{10}}{k_1}x_5\right\}, \tag{6}$$

where $D_1 = 2a_1x_5 + b_1$. With $V = \frac{1}{2}(e_1^2 + e_2^2 + e_3^2)$, the closed-loop error dynamics of the design model are $\dot{e}_1 = -k_{p1}e_1 + e_2$, $\dot{e}_2 = -e_1 - k_{p2}e_2 + e_3$, $\dot{e}_3 = -e_2 - k_{p3}e_3$, so that $\dot{V} = -k_{p1}e_1^2 - k_{p2}e_2^2 - k_{p3}e_3^2$ in the absence of saturation and coupling. The yaw law U_2 follows the same three steps on $\dot{x}_3 = x_4$, $I_2\dot{x}_4 = G_2 := F_2(x_6) - B_{1\phi}x_4 - B_{2\phi}x_3$, $\dot{x}_6 = (-T_{20}x_6 + k_2U_2)/T_{21}$, with errors e_4, e_5, e_6, gains k_{y1}, k_{y2}, k_{y3} and $D_2 = 2a_2x_6 + b_2$. In both laws the denominators D_1 and D_2 are kept away from zero by a guard $|D_i| \geq 0.02$, and the true states are replaced by the observer estimates. The gains used throughout are $(k_{p1}, k_{p2}, k_{p3}) = (7, 5, 2)$ and $(k_{y1}, k_{y2}, k_{y3}) = (10, 10, 5)$. Although the two laws are derived separately, they are evaluated and applied simultaneously at every control step, so that both axes are controlled at the same time. The coupling terms omitted from the design models, namely the gyroscopic moment on the pitch axis and the main-rotor reaction torque on the yaw axis, are therefore physically present during operation and act on each loop as disturbances generated by the motion of the other axis. The observer, in contrast, is built on the full coupled model and accounts for these terms explicitly.

Remark 1. *Since $|U_1| \leq \bar{U} = 2.5$ V and $\dot{x}_5 = (-T_{10}x_5 + k_1U_1)/T_{11}$ is a stable first-order lag, any torque trajectory that starts with $|x_5(0)| \leq k_1\bar{U}/T_{10} = 2.75$ satisfies $|x_5(t)| \leq 2.75$ for all t. Hence $D_1 = 2a_1x_5 + b_1 \in [0.0182, 0.1667]$, which is strictly positive: the quadratic map F_1 remains monotone on the reachable set, whose lower end lies well above its vertex at $x_5 = -b_1/(2a_1) = -3.42$. The lower bound is reached only after a sustained saturation $U_1 = -2.5$ V lasting several time constants of the main motor. On the rig, the longest uninterrupted episode at $U_1 = -2.5$ V lasted 0.95, 0.92 and 1.14 s in the three experiments of Section 3.2, that is, about one time constant of the main motor. The estimated torque used in (6) and in the observer can transiently leave the reachable set, which is why the guard is kept; along the recorded trajectories it remained above -1.91, so that $D_1 \geq 0.041$ and the guard $|D_1| \geq 0.02$ was never active, and the condition number of the observer Jacobian $\partial\Phi/\partial\mathbf{x}$ in (8) remained below 55 (median 28). The near-singular configuration noted for sustained negative saturation is therefore excluded on the reachable set and was not approached in the experiments. A signed thrust map $a_1x_5|x_5| + b_1x_5$, for which $D_1 = 2a_1|x_5| + b_1 \geq b_1$, would remove the guard altogether and is discussed in Section 4.*

2.3. High-Gain Observer with Saturated Injection

The observer is built on the full coupled model (1), written as $\dot{\mathbf{x}} = f(\mathbf{x}) + g\,U$. For each output $h_1 = x_1$ and $h_2 = x_3$ the relative degree is three, and the observability map

$$\Phi(\mathbf{x}) = \big[h_1,\ L_fh_1,\ L_f^2h_1,\ h_2,\ L_fh_2,\ L_f^2h_2\big]^T = \big[x_1,\ x_2,\ \dot{x}_2,\ x_3,\ x_4,\ \dot{x}_4\big]^T \tag{7}$$

is a diffeomorphism on the operating domain. The observer reads [12,13]

$$\dot{\hat{\mathbf{x}}} = f(\hat{\mathbf{x}}) + g\,U + \left[\frac{\partial\Phi}{\partial\mathbf{x}}(\hat{\mathbf{x}})\right]^{-1}\Theta\,\Delta^{-1}(\boldsymbol{\theta})\,K\,\mathrm{sat}_{\bar{\varepsilon}}\left(y - C\hat{\mathbf{x}}\right), \tag{8}$$

where $\Delta(\boldsymbol{\theta}) = \text{blockdiag}\left(\Delta_3(\theta_\psi), \Delta_3(\theta_\phi)\right)$ with $\Delta_3(\theta) = \text{diag}(1, \theta^{-1}, \theta^{-2})$, $\Theta = \text{blockdiag}(\theta_\psi I_3, \theta_\phi I_3)$, $K = \text{blockdiag}(\bar{k}, \bar{k})$ with $\bar{k} = [3, 3, 1]^T$, and $\text{sat}_{\bar{\varepsilon}}$ is the componentwise saturation at $\bar{\varepsilon} = 0.05$ rad. With this scaling the injection vector of each axis is $\theta \Delta_3^{-1}(\theta)\bar{k}\,\varepsilon = [3\theta,\ 3\theta^2,\ \theta^3]^T\varepsilon$, and the linearised error dynamics of each axis have the characteristic polynomial $(s+\theta)^3$ for every value of θ. The nonzero entries of the Jacobian that are not trivially equal to one are

$$\begin{aligned}
&\frac{\partial \dot{x}_2}{\partial x_1} = \frac{-M_g \cos x_1 + K_{gy} x_4 \sin x_1\, F_1}{I_1}, && \frac{\partial \dot{x}_2}{\partial x_2} = -\frac{B_{1\psi}}{I_1},\\
&\frac{\partial \dot{x}_2}{\partial x_4} = -\frac{K_{gy} \cos x_1\, F_1}{I_1}, && \frac{\partial \dot{x}_2}{\partial x_5} = \frac{D_1\,(1 - K_{gy} x_4 \cos x_1)}{I_1},\\
&\frac{\partial \dot{x}_4}{\partial x_4} = -\frac{B_{1\phi}}{I_2}, && \frac{\partial \dot{x}_4}{\partial x_5} = -\frac{1.75\,K_c\,D_1}{I_2},\\
&\frac{\partial \dot{x}_4}{\partial x_6} = \frac{D_2}{I_2}.
\end{aligned} \tag{9}$$

The saturated injection limits the peaking that a large initial error or an encoder glitch would otherwise produce [11,19], and the guards on D_1 and D_2 keep the Jacobian invertible.

Remark 2. *A frequent implementation pitfall is to scale the gain vector twice, using $[3\theta, 3\theta^2, \theta^3]^T$ together with $\Delta^{-1}(\theta)$. The resulting polynomial $s^3 + 3\theta s^2 + 3\theta^3 s + \theta^5$ is Hurwitz only for $\theta < 9$, so that any scheduler allowed to exceed this value can destabilise the observer. The scaling in (8) removes this restriction and gives θ the meaning of an observer bandwidth.*

2.4. Practical Stability with a Time-Varying Gain

Let $z = \Phi(\mathbf{x})$, $\hat{z} = \Phi(\hat{\mathbf{x}})$ and, for each axis $a \in \{\psi, \phi\}$, $\eta_a = \Delta_3(\theta_a)(z_a - \hat{z}_a)$, with $\eta = [\eta_\psi^T, \eta_\phi^T]^T$. In observability coordinates the model reads $\dot{z} = A_6 z + B_6 \varphi(z, U)$, where $A_6 = \text{blockdiag}(A, A)$, A is the 3×3 shift matrix, $B_6 = \text{blockdiag}(e_3, e_3)$ and φ collects the third derivatives of the two outputs. Inside the region where the injection is not saturated, standard computations [11,16] give

$$\dot{\eta} = M(t)\,\eta + \sum_{a \in \{\psi,\phi\}} \frac{\dot{\theta}_a}{\theta_a} D_a\,\eta + w(t), \qquad M(t) = \text{blockdiag}\left(\theta_\psi \bar{A},\ \theta_\phi \bar{A}\right) + \Gamma(t), \tag{10}$$

where $\bar{A} = A - \bar{k} C_0$, $C_0 = [1, 0, 0]$, $D_\psi = \text{blockdiag}(D, 0_3)$ and $D_\phi = \text{blockdiag}(0_3, D)$ with $D = \text{diag}(0, -1, -2)$, $\Gamma(t) = \Delta(\boldsymbol{\theta}) B_6\, \bar{\varphi}_z(t)\, \Delta^{-1}(\boldsymbol{\theta})$ with $\bar{\varphi}_z$ the mean-value Jacobian of φ between $\hat{z}$ and z, and w gathers the model mismatch and the measurement perturbation ν, entering as $\Delta B_6 \delta - \text{blockdiag}(\theta_\psi \bar{k}, \theta_\phi \bar{k})\,\nu$.

Assumption 1. *The closed-loop trajectories remain in a compact set on which $\|w(t)\| \le \bar{w}$, $\theta_a \in [\theta_{\min,a}, \theta_{\max,a}]$, and θ_a is absolutely continuous.*

Proposition 1. *Let $P = \text{blockdiag}(P_\psi, P_\phi) \succ 0$ and let $\sigma(t)$ satisfy $PM(t) + M(t)^T P \preceq -2\sigma(t) P$ along the trajectory. Define $\mu = \max_a \left\| P_a^{-1/2} (P_a D + D P_a) P_a^{-1/2} \right\|$, $\rho(t) = \max_a |\dot{\theta}_a(t)| / \theta_a(t)$ and $\kappa = \lambda_{\max}(P)/\lambda_{\min}(P)$. Under Assumption 1, if there exist $c > 0$ and $N_0 \ge 0$ such that*

$$\int_s^t \left(\sigma(\tau) - \tfrac{1}{2}\mu\rho(\tau)\right) d\tau \ \ge\ c\,(t - s) - N_0 \qquad \textit{for all } t \ge s \ge t_0, \tag{11}$$

then

$$\|\eta(t)\| \le \sqrt{\kappa}\, e^{N_0} \left(e^{-c(t - t_0)} \|\eta(t_0)\| + \frac{\bar{w}}{c} \right), \tag{12}$$

that is, the scaled estimation error is uniformly ultimately bounded with an ultimate bound $\sqrt{\kappa}\, e^{N_0} \bar{w}/c$.

Proof. Let $V = \eta^T P \eta$ and $W = \sqrt{V}$. Since P is block diagonal, $|\eta^T (P D_a + D_a P)\eta| \leq \mu\, \eta_a^T P_a \eta_a$, so the sum in (10) contributes at most $\mu\rho V$ to $\dot{V}$. Hence $\dot{V} \leq -(2\sigma - \mu\rho)V + 2\sqrt{V}\, \|P^{1/2} w\|$ and $\dot{W} \leq -(\sigma - \mu\rho/2)W + \sqrt{\lambda_{\max}(P)}\, \bar{w}$. The comparison lemma and (11) give $W(t) \leq e^{N_0} e^{-c(t-t_0)} W(t_0) + e^{N_0} \sqrt{\lambda_{\max}(P)}\, \bar{w}/c$, and $\lambda_{\min}(P)\|\eta\|^2 \leq V \leq \lambda_{\max}(P)\|\eta\|^2$ yields the bound. □

Condition (11) only requires the decay rate to dominate the gain variation on average, which suits a scheduler that changes the gain in short bursts. If $\sigma(t) \geq \sigma_0$ and $\rho(t) \leq \bar{\rho}$ for all t, it holds with $N_0 = 0$ whenever $\bar{\rho} < 2\sigma_0/\mu$. Two evaluations of the constants are given below. Since $\eta = \Delta(\boldsymbol{\theta})(z - \hat{z})$, the bound on the torque-related estimation error is further multiplied by θ^2, which is why encoder quantisation appears in the control voltage, and the noise part of $\bar{w}$ grows with $\theta_{\max}$ whereas the decay rate grows with $\theta_{\min}$. Combined with the input-to-state property of the backstepping error dynamics with respect to the estimation error, the proposition gives practical stability of the output-feedback loop in the sense of [11].

Remark 3 (Nominal evaluation). *If the nonlinear term is neglected,* $\Gamma \equiv 0$, *and* $P_a \bar{A} + \bar{A}^T P_a \preceq -\lambda_0 P_a$ *gives* $2\sigma(t) \geq \lambda_0 \theta_{\min}$. *The ratio* λ_0/μ *is maximised over* $P_a \succ 0$ *by a generalised eigenvalue problem, which yields* $\lambda_0 = 1.233$, $\mu = 5.835$ *and* $\lambda_0/\mu = 0.211$. *The pointwise rate bound is then* $\bar{\rho} < 0.211\, \theta_{\min}$, *that is* $0.85\ s^{-1}$ *for pitch (*$\theta_{\min} = 4$*) and* $0.63\ s^{-1}$ *for yaw (*$\theta_{\min} = 3$*). With the solution of the Lyapunov equation* $P\bar{A} + \bar{A}^T P = -I_3$ *and the bound* $\|PD + DP\| = 16.39$, *the same reasoning only gives* $\bar{\rho} < 0.24\ s^{-1}$ *for* $\theta_{\min} = 4$, *about 3.5 times less.*

Remark 4 (Evaluation along the recorded operating points). *The norm of* $\Gamma(t)$, *evaluated at the estimated state along the three experiments of Section 3.2 (6000 points sampled every 40 ms for* $t \geq 20$ *s), has a median of 6.0 and a maximum of 8.7, so that a bound of the form* $\|\Gamma\| \leq L$ *cannot be dominated by* $\theta_{\min}\bar{A}$ *for* $\theta_{\min} \in \{3, 4\}$. *The dominant entries of* Γ *are, however, damping terms of the rotor and friction dynamics, of the order of* -1 *for pitch and* -6 *for yaw, which are stabilising. Exploiting this structure, every frozen matrix* $M(t)$ *at the sampled points is Hurwitz, with a largest real part of* -2.25. *A block-diagonal* P *computed by a linear matrix inequality [36] on a subset of these points, and verified on all of them, gives* $2\sigma(t) \geq 2.60$ *(mean 4.96) with* $\mu = 5.78$. *The pointwise condition therefore holds on the recorded operating envelope for any scheduler with* $\bar{\rho} < 0.45\ s^{-1}$. *Table 3 evaluates the averaged condition (11) on the recorded gain trajectories. The three schedulers change the gain faster than this sufficient condition allows: the condition holds on average for ML2 only, and over 10 s windows for none of them. Proposition 1 therefore does not certify the recorded experiments, whose stability remains an experimental observation. It does certify a rate-limited variant of the scheduler with symmetric rate limits of* $0.4\ s^{-1}$, *whose performance is evaluated in Section 3.3.*

Table 3. Evaluation of the averaged condition (11) along the recorded experiments ($t \geq 20$ s, $\mu = 5.78$). The condition requires $\bar{\rho} < 2\bar{\sigma}/\mu$ on average and over every window.

Experiment	Mean 2σ (s^{-1})	Threshold $2\bar{\sigma}/\mu$ (s^{-1})	Mean ρ (s^{-1})	10 s Windows Satisfying (11) (%)
IT3-CA	5.28	0.91	1.37	0
ML1	4.80	0.83	1.03	2
ML2	4.81	0.83	0.79	64

2.5. Feature Extraction

All three schedulers run at $T_s = 0.01$ s and use the same three features per axis, computed from the output-estimation error $\varepsilon_k = y_k - \hat{y}_k$ of that axis and its control voltage U_k.

2.5.1. Noise Floor

The increments of ε are dominated by sensor noise, because the estimated output varies by much less than one encoder count per sample. For white noise of standard deviation σ_n, $\mathbb{E}|\varepsilon_k - \varepsilon_{k-1}| = 2\sigma_n/\sqrt{\pi} \approx 1.13\,\sigma_n$. The noise floor is therefore estimated as

$$\hat{n}_k = (1-g)\hat{n}_{k-1} + g|\varepsilon_k - \varepsilon_{k-1}|, \qquad \varepsilon_{0,k} = \max\left(n_q\, q,\ 2.5\,\hat{n}_k/1.13\right), \tag{13}$$

with $g = T_s/(0.5 + T_s)$, the encoder resolution $q = 2\pi/2048$ rad and $n_q = 0.3$.

2.5.2. Excess Error and Trend

With $\bar{e}_k = (1-a)\bar{e}_{k-1} + a|\varepsilon_k|$, $a = T_s/(0.05 + T_s)$, the first two features are

$$\begin{aligned} x_{1,k} &= \mathrm{sat}_{[0,1]}\left(\frac{\bar{e}_k - \varepsilon_{0,k}}{e_s}\right), \\ x_{2,k} &= \mathrm{sat}_{[-1,1]}\left(\frac{[\bar{e}_k - \varepsilon_{0,k}]_+ - [\bar{e}_{k-1} - \varepsilon_{0,k-1}]_+}{T_s\, e_s/0.1}\right), \end{aligned} \tag{14}$$

with $e_s = 0.008$ rad and $[\cdot]_+ = \max(\cdot, 0)$. Only the part of the error that exceeds the noise floor contributes to the trend, so that quantisation alone does not trigger a gain increase.

2.5.3. Chattering Index

The second difference of the control voltage removes the smooth part of U that follows the reference and retains its high-frequency content:

$$c_k = (1-b)c_{k-1} + b\,|U_k - 2U_{k-1} + U_{k-2}|, \qquad x_{3,k} = \mathrm{sat}_{[0,1]}(c_k/c_{\mathrm{ref}}), \tag{15}$$

with $b = T_s/(0.2 + T_s)$ and $c_{\mathrm{ref}} = 0.1$ V. The second difference is taken over the scheduler period T_s, so that c_{ref} is calibrated for an execution every 10 ms. All three schedulers are executed at this period in the experiments reported here.

2.6. Chattering-Aware Interval Type-3 Fuzzy Scheduler

2.6.1. IT3 Sets

Each input x_i is partitioned by Gaussian sets $G(x; c, s) = \exp\big(-(x-c)^2/(2s^2)\big)$ with centres $\{0, 0.5, 1\}$ for x_1 (Small, Medium, Large), $\{-1, 0, 1\}$ for x_2 (Decreasing, Zero, Increasing) and $\{0, 1\}$ for x_3 (Low, High). The base IT2 set of input i has an upper membership of width $\bar{s}_i$ and a lower membership of width $\underline{s}_i < \bar{s}_i$. In the IT3 set both widths are themselves uncertain with a spread δ_i, and the vertical slices at level α are delimited by four Gaussian curves (Figure 2):

$$\begin{aligned} \bar{\mu}^{\mathrm{o}}_{\alpha} &= G\big(x; c, \bar{s}_i + (1-\alpha)\delta_i\big), & \underline{\mu}^{\mathrm{o}}_{\alpha} &= G\big(x; c, \underline{s}_i + (1-\alpha)\delta_i\big), \\ \bar{\mu}^{\mathrm{i}}_{\alpha} &= G\big(x; c, \bar{s}_i - (1-\alpha)\delta_i\big), & \underline{\mu}^{\mathrm{i}}_{\alpha} &= G\big(x; c, \underline{s}_i - (1-\alpha)\delta_i\big), \end{aligned} \tag{16}$$

which define an outer IT2 set $[\underline{\mu}^{\mathrm{o}}_{\alpha}, \bar{\mu}^{\mathrm{o}}_{\alpha}]$ and an inner IT2 set $[\underline{\mu}^{\mathrm{i}}_{\alpha}, \bar{\mu}^{\mathrm{i}}_{\alpha}]$. Both collapse onto the base IT2 set at $\alpha = 1$. The parameters are $\bar{s} = [0.28, 0.55, 0.55]$, $\underline{s} = [0.18, 0.35, 0.35]$ and $\delta = [0.04, 0.08, 0.08]$.

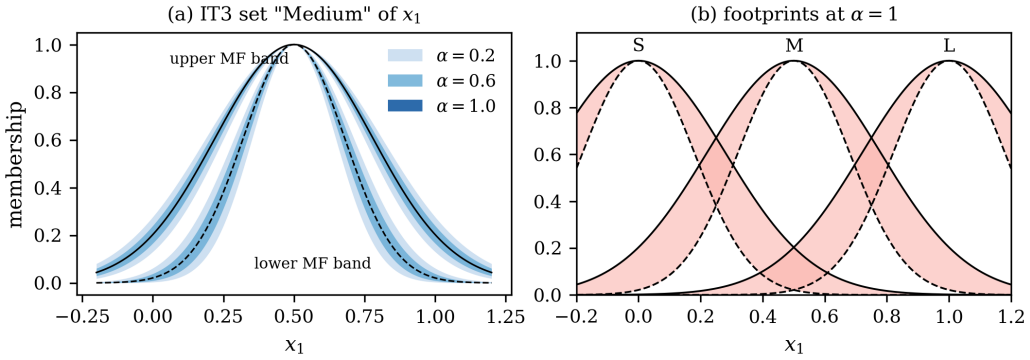

Figure 2. Interval type-3 sets of the first input. (**a**) Set Medium: the bands around the upper (solid) and lower (dashed) membership functions represent the third-dimension uncertainty at $\alpha = 0.2$ and $\alpha = 0.6$; at $\alpha = 1$ they collapse onto the two curves. (**b**) Footprints of uncertainty of the three sets at $\alpha = 1$.

2.6.2. Rule Base

The 18 rules have singleton consequents on the normalised gain $\Psi \in [0, 1]$, with the levels VL $= 0$, L $= 0.2$, M $= 0.45$, H $= 0.7$ and VH $= 1$ (Table 4). When the chattering is low, the gain rises with the excess error and its trend. When the chattering is high, every consequent is shifted down by one or two levels. The combination (Small, Zero) keeps a low gain in both cases, which reproduces the damped-zero principle of [20] and prevents the observer from chasing quantisation noise in steady state.

Table 4. IT3-CA rule base: consequent level of Ψ for each combination of excess error x_1, trend x_2 and chattering x_3.

x_1	x_3 Low			x_3 High		
	D	**Z**	**I**	**D**	**Z**	**I**
Small	VL	L	M	VL	VL	L
Medium	L	M	H	VL	L	M
Large	M	H	VH	L	M	H

2.6.3. Inference and Type Reduction

Firing strengths are obtained with the product t-norm on each of the four membership curves, giving for rule r at level α the outer interval $[\underline{f}^{\mathrm{o}}_r, \bar{f}^{\mathrm{o}}_r]$ and the inner interval $[\underline{f}^{\mathrm{i}}_r, \bar{f}^{\mathrm{i}}_r]$. The Karnik–Mendel algorithm [29] is applied to both IT2 slices with the rules sorted by consequent, giving four end points $y^{\mathrm{o}}_{l,\alpha}$, $y^{\mathrm{o}}_{r,\alpha}$, $y^{\mathrm{i}}_{l,\alpha}$ and $y^{\mathrm{i}}_{r,\alpha}$. The output of each alpha-plane and the crisp output are

$$y_\alpha = \tfrac{1}{4}\left(y^{\mathrm{o}}_{l,\alpha} + y^{\mathrm{o}}_{r,\alpha} + y^{\mathrm{i}}_{l,\alpha} + y^{\mathrm{i}}_{r,\alpha}\right), \qquad \Psi = \frac{\sum_{\alpha\in\mathcal{A}} \alpha\, y_\alpha}{\sum_{\alpha\in\mathcal{A}} \alpha}, \tag{17}$$

with $\mathcal{A} = \{0.2,\ 0.6,\ 1\}$, following the alpha-plane representation of higher-type sets [22,27,28].

Remark 5 (Role of the third dimension). *The outer and inner slices widen and narrow the base widths by the same amount* $(1 - \alpha)\delta_i$. *The four end points in (17) are piecewise smooth functions of the widths, so that their first-order variations with respect to* $(1 - \alpha)\delta_i$ *cancel in the average* y_α, *which differs from the output of the base IT2 system by a term of order* δ_i^2. *With the spreads used here,* $\delta_i/\bar{s}_i \le 0.15$, *this term is small. The IT2 variant obtained with* $\mathcal{A} = \{1\}$ *is evaluated in Section 3.3. The IT3 formulation is retained as the general framework, because it allows an*

asymmetric spread or a different alpha-plane weighting without changing the scheduler structure, but in the present configuration it behaves as its IT2 core.

2.6.4. Gain Mapping and Rate Limiter

The target gain $\theta_k^\star = \theta_{\min} + (\theta_{\max} - \theta_{\min})\Psi_k$ is reached through an asymmetric relative rate limiter,

$$\theta_{k+1} = \theta_k + \mathrm{sat}\left(\theta_k^\star - \theta_k;\ -\rho_{\mathrm{dn}}\theta_k T_s,\ \rho_{\mathrm{up}}\theta_k T_s\right), \tag{18}$$

with $\rho_{\mathrm{up}} = 5\ \mathrm{s}^{-1}$ and $\rho_{\mathrm{dn}} = 1\ \mathrm{s}^{-1}$, so that the gain rises quickly during transients and decreases slowly afterwards. The ranges are $[\theta_{\min}, \theta_{\max}] = [4, 25]$ for pitch and $[3, 20]$ for yaw. The limiter enforces the bounds of Assumption 1 by construction. One update of the scheduler for one axis is summarised in Algorithm 1.

Algorithm 1 IT3-CA update of the observer gain for one axis

Require: ε_k, U_k; stored ε_{k-1}, U_{k-1}, U_{k-2}, $\hat{n}_{k-1}$, $\bar{e}_{k-1}$, c_{k-1}, θ_k

Ensure: θ_{k+1}

$\hat{n}_k \leftarrow (1-g)\hat{n}_{k-1} + g|\varepsilon_k - \varepsilon_{k-1}|$; $\quad \varepsilon_{0,k} \leftarrow \max(n_q q,\ 2.5\hat{n}_k/1.13)$

$\bar{e}_k \leftarrow (1-a)\bar{e}_{k-1} + a|\varepsilon_k|$; $\quad c_k \leftarrow (1-b)c_{k-1} + b|U_k - 2U_{k-1} + U_{k-2}|$

compute the features $x_{1,k}$, $x_{2,k}$ and $x_{3,k}$

for $\alpha \in \{0.2,\ 0.6,\ 1\}$ **do**

    evaluate the outer and inner upper and lower Gaussian bounds of the eight input sets

    compute the outer and inner firing intervals of the 18 rules (product t-norm)

    apply Karnik–Mendel to the outer and inner IT2 slices to obtain four end points

    $y_\alpha \leftarrow \frac{1}{4}\left(y^{\mathrm{o}}_{l,\alpha} + y^{\mathrm{o}}_{r,\alpha} + y^{\mathrm{i}}_{l,\alpha} + y^{\mathrm{i}}_{r,\alpha}\right)$

end for

$\Psi_k \leftarrow \sum_\alpha \alpha\, y_\alpha \big/ \sum_\alpha \alpha$; $\quad \theta_k^\star \leftarrow \theta_{\min} + (\theta_{\max} - \theta_{\min})\Psi_k$

$\theta_{k+1} \leftarrow \theta_k + \mathrm{sat}\left(\theta_k^\star - \theta_k;\ -\rho_{\mathrm{dn}}\theta_k T_s,\ \rho_{\mathrm{up}}\theta_k T_s\right)$

2.7. Neural Schedulers Used as Benchmarks

Two neural schedulers are designed as benchmarks. They receive the same feature vector $\mathbf{x}_k = [x_{1,k},\ x_{2,k},\ x_{3,k}]^T$ as IT3-CA, they produce the normalised gain $\Psi_k \in (0,1)$ through a logistic output $\Psi = \sigma_l(z) = 1/(1+e^{-z})$, and they use the same gain mapping and rate limiter (18). Their difference with IT3-CA lies in the way the map from $\mathbf{x}_k$ to Ψ_k is obtained: IT3-CA encodes it in a fixed rule base, whereas the neural schedulers learn it online.

2.7.1. Performance Cost

Our earlier neural tuner [21] was trained towards a fixed analytic target, so that the network only reproduced a predefined function of the error and received no information on the closed-loop behaviour. Here both networks descend instead the instantaneous performance cost

$$J_k = \tfrac{1}{2}\left(w_e\, x_{1,k}^2 + w_c\, x_{3,k}^2 + w_0\left(\frac{\theta_k - \theta_0}{\theta_0}\right)^2\right), \tag{19}$$

which penalises the excess estimation error, the chattering of the control voltage and the distance to a nominal gain θ_0. The last term regularises the gain when both features vanish. The gradient of J with respect to θ requires the sensitivities of the features, which are not available in closed form. They are approximated by two scaling laws suggested by Proposition 1: the model-mismatch part of the estimation error decreases as $1/\theta$, which gives $\partial x_1/\partial\theta \approx -x_1/\theta$, and the quantisation noise injected into the torque estimate grows

as θ^2, which gives $\partial x_3/\partial\theta \approx k_c\, x_3/\theta$ with $k_c = 2$. Since $\theta = \theta_{\min} + (\theta_{\max} - \theta_{\min})\sigma_l(z)$ and $\sigma_l'(z) = \Psi(1-\Psi)$, the error signal at the network output is

$$\delta_z := \frac{\partial J}{\partial z} = \Big(\frac{-w_e\, x_1^2 + w_c\, k_c\, x_3^2}{\theta} + w_0 \frac{\theta - \theta_0}{\theta_0^2}\Big)\,(\theta_{\max} - \theta_{\min})\,\Psi(1-\Psi). \tag{20}$$

The weights are $w_e = 1$, $w_c = 0.5$ and $w_0 = 4$, and the nominal gains are $\theta_0 = 7$ for pitch and $\theta_0 = 6$ for yaw. The sign structure of (20) is the key property: an excess estimation error pushes the gain up, chattering pushes it down, and the nominal term restores θ_0 when the measurement is clean.

2.7.2. MLP-ADAM Scheduler (ML1)

ML1 is a 3-6-1 multilayer perceptron with hyperbolic tangent hidden units,

$$\mathbf{h} = \tanh\big(W_1\mathbf{x} + \mathbf{b}_1\big), \qquad z = W_2\mathbf{h} + b_2, \qquad \Psi = \sigma_l(z), \tag{21}$$

with $W_1 \in \mathbb{R}^{6\times 3}$, $\mathbf{b}_1 \in \mathbb{R}^6$, $W_2 \in \mathbb{R}^{1\times 6}$ and $b_2 \in \mathbb{R}$, that is, 31 trainable parameters. Back-propagating δ_z gives the hidden error $\boldsymbol{\delta}_h = \delta_z\, W_2^T \odot (\mathbf{1} - \mathbf{h}\odot\mathbf{h})$ and the gradients

$$\frac{\partial J}{\partial W_2} = \delta_z\, \mathbf{h}^T, \qquad \frac{\partial J}{\partial b_2} = \delta_z, \qquad \frac{\partial J}{\partial W_1} = \boldsymbol{\delta}_h\, \mathbf{x}^T, \qquad \frac{\partial J}{\partial \mathbf{b}_1} = \boldsymbol{\delta}_h. \tag{22}$$

Each parameter p with gradient g_k is updated with ADAM [37],

$$m_k = \beta_1 m_{k-1} + (1-\beta_1) g_k, \quad v_k = \beta_2 v_{k-1} + (1-\beta_2) g_k^2, \quad p \leftarrow p - \eta_a \frac{m_k/(1-\beta_1^k)}{\sqrt{v_k/(1-\beta_2^k)} + \epsilon}, \tag{23}$$

with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\eta_a = 2\times 10^{-3}$ and $\epsilon = 10^{-8}$. ADAM normalises each step by a running estimate of the gradient magnitude, which makes the learning rate insensitive to the scale of δ_z. The initialisation is deterministic, $[W_1]_{ij} = 0.5\sin(1.3\,ij)$, $[W_2]_i = 0.3\cos(0.7\,i)$, $\mathbf{b}_1 = \mathbf{0}$ and $b_2 = \ln\big(\Psi_0/(1-\Psi_0)\big)$ with $\Psi_0 = (\theta_0 - \theta_{\min})/(\theta_{\max} - \theta_{\min})$, so that the scheduler starts at the nominal gain and every run is reproducible. Algorithm 2 summarises one update.

Algorithm 2 ML1 (MLP-ADAM) update of the observer gain for one axis

Require: ε_k, U_k; network parameters, ADAM moments, θ_k

Ensure: θ_{k+1}

1: compute $\mathbf{x}_k$ as in Algorithm 1, lines 1 to 3
2: $\mathbf{h} \leftarrow \tanh(W_1\mathbf{x}_k + \mathbf{b}_1)$; $\quad z \leftarrow W_2\mathbf{h} + b_2$; $\quad \Psi_k \leftarrow \sigma_l(z)$
3: $\delta_z \leftarrow$ (20) evaluated at $\mathbf{x}_k$, θ_k and Ψ_k
4: $\boldsymbol{\delta}_h \leftarrow \delta_z\, W_2^T \odot (\mathbf{1} - \mathbf{h}\odot\mathbf{h})$
5: form the gradients $\delta_z\mathbf{h}^T$, δ_z, $\boldsymbol{\delta}_h\mathbf{x}_k^T$, $\boldsymbol{\delta}_h$
6: update W_2, b_2, W_1, $\mathbf{b}_1$ with ADAM
7: $\theta_k^\star \leftarrow \theta_{\min} + (\theta_{\max} - \theta_{\min})\Psi_k$ and apply the rate limiter (18)

2.7.3. RBF-NLMS Scheduler (ML2)

ML2 is a radial basis function network [38] with $M = 18$ Gaussian units whose centres $\mathbf{c}_j$ form the same $3\times 3\times 2$ grid as the fuzzy partition, $\{0, 0.5, 1\}\times\{-1, 0, 1\}\times\{0, 1\}$:

$$\varphi_j(\mathbf{x}) = \exp\Big(-\frac{1}{2}\sum_{i=1}^{3}\frac{(x_i - c_{j,i})^2}{s_i^2}\Big), \qquad z = \mathbf{w}^T\boldsymbol{\varphi}(\mathbf{x}) + b, \qquad \Psi = \sigma_l(z), \tag{24}$$

with widths $\mathbf{s} = [0.3, 0.6, 0.6]$. The centres and widths are fixed, and the 19 output parameters $(\mathbf{w}, b)$ are trained with the normalised least-mean-squares rule [39],

$$\mathbf{w} \leftarrow \mathbf{w} - \eta \frac{\delta_z \boldsymbol{\varphi}}{\boldsymbol{\varphi}^T \boldsymbol{\varphi} + 1}, \qquad b \leftarrow b - \eta \frac{\delta_z}{\boldsymbol{\varphi}^T \boldsymbol{\varphi} + 1}, \tag{25}$$

with $\eta = 0.05$ and $\mathbf{w}$ initialised at zero and $b = \ln\left(\Psi_0/(1-\Psi_0)\right)$. Because the output is linear in $(\mathbf{w}, b)$, this rule is a normalised gradient step whose effective step size does not depend on the activation level of the network, and it avoids the inversion of an ill-conditioned matrix. The localised receptive fields give ML2 a structure close to a zero-order fuzzy system with learned consequents, which makes it the neural counterpart closest to IT3-CA. Algorithm 3 summarises one update.

Algorithm 3 ML2 (RBF-NLMS) update of the observer gain for one axis

Require: ε_k, U_k; output weights $\mathbf{w}$, bias b, θ_k

Ensure: θ_{k+1}

1: compute $\mathbf{x}_k$ as in Algorithm 1, lines 1 to 3
2: **for** $j = 1, \ldots, 18$ **do**
3: $\quad \varphi_j \leftarrow \exp\left(-\frac{1}{2}\sum_i (x_{i,k} - c_{j,i})^2 / s_i^2\right)$
4: **end for**
5: $z \leftarrow \mathbf{w}^T \boldsymbol{\varphi} + b; \quad \Psi_k \leftarrow \sigma_l(z)$
6: $\delta_z \leftarrow$ (20) evaluated at $\mathbf{x}_k$, θ_k and Ψ_k
7: $\mathbf{w} \leftarrow \mathbf{w} - \eta\, \delta_z \boldsymbol{\varphi}/(\boldsymbol{\varphi}^T \boldsymbol{\varphi} + 1); \quad b \leftarrow b - \eta\, \delta_z/(\boldsymbol{\varphi}^T \boldsymbol{\varphi} + 1)$
8: $\theta_k^\star \leftarrow \theta_{\min} + (\theta_{\max} - \theta_{\min})\Psi_k$ and apply the rate limiter (18)

2.7.4. Comparison of the Three Schedulers

Table 5 contrasts the three designs. IT3-CA has no online learning: its 35 design parameters are fixed, and its adaptation comes from the nonlinear map encoded in the rule base. The neural schedulers learn online with 31 and 19 trainable parameters. IT3-CA is the most demanding in computation because of its 96 Gaussian evaluations and 12 Karnik–Mendel procedures per update, each limited to ten iterations, so that its computation is bounded; the measured execution times are reported in Section 3.4.

Table 5. Structural comparison of the three schedulers (per axis and per update).

	IT3-CA	**ML1 (MLP-ADAM)**	**ML2 (RBF-NLMS)**
Inputs	x_1, x_2, x_3	x_1, x_2, x_3	x_1, x_2, x_3
Map from inputs to Ψ	18 IT3 rules, 3 alpha-planes	3-6-1 perceptron	18 Gaussian units
Online learning	none	ADAM on cost (19)	NLMS on cost (19)
Parameters	35 fixed	31 trainable	19 trainable, 57 fixed
Nonlinear evaluations	96 exponentials	6 tanh, 1 exponential, 31 square roots	19 exponentials
Iterative procedures	12 Karnik–Mendel	none	none

The complete closed loop is shown in Figure 3. The three schedulers differ only in the block that maps the features to the observer gains.

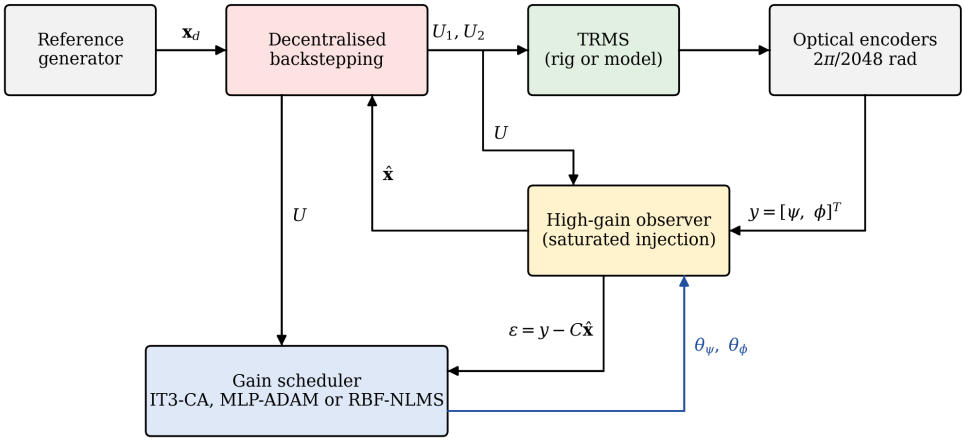

Figure 3. Closed-loop architecture. The scheduler block is the only element that differs between the three configurations compared in this study. Black arrows show the signal flow of the control loop, blue elements show the scheduling path that sets the observer gains θ_ψ and θ_ϕ, and the block colours distinguish the controller, the plant, the observer and the scheduler.

2.8. Tuning Procedure

The design parameters of the three schedulers were fixed before the experiments reported here and were not adjusted to the recorded data. The shared elements, namely the features, the gain ranges, the rate limiter and the update period, are identical by construction. For IT3-CA, the consequent levels were laid out as a monotone grid in the excess error and its trend, the High-chattering branch was shifted down by one or two levels, and c_{ref} was set to the order of magnitude of the second difference of the control voltage over one scheduler period at the gain floor. For the neural schedulers, w_e and w_c express the same priority between estimation and chattering as the rule base, w_0 sets the strength of the return to the nominal gain, and θ_0 was chosen close to the mean gain of a fixed-gain observer that tracks the references without visible chattering. The learning rates were the largest values that gave a monotone decrease in the cost of simulation. In each family, one parameter governs the position on the accuracy and smoothness trade-off: c_{ref} for IT3-CA and w_c for the neural schedulers. Both are swept in Section 3.3, so that the families are compared through their trade-off curves rather than through a single tuning point.

2.9. Experimental Setup and Performance Metrics

The experiments were carried out on a Feedback Instruments TRMS 33-220 interfaced with a host computer through an Advantech PCI-1711 data acquisition card [2]. The control loop was compiled with the Real-Time Windows Target of MATLAB/Simulink R2014a and executed at 1 kHz with the fixed-step Euler solver. All three schedulers were executed every 10 ms, that is, once every ten control steps, with the coefficients specified for $T_s = 0.01$ s, so that their filter time constants, rate limits, chattering calibration and learning periods coincide. The pitch and yaw schedulers used the ranges and nominal gains given above. The optical encoders provide 2048 counts per revolution, so that the measured angles are quantised with a step of $q = 3.07$ mrad. In every experiment both axes were free and were driven simultaneously, that is, the TRMS operated with its two degrees of freedom and its full cross-coupling. The references $\psi_d = 0.4\sin(0.5t)$ rad and $\phi_d = 0.8\sin(0.5t)$ rad were applied together over 100 s. All three configurations share the controller gains, the observer, the features and the saturation limits given above. Table 6 summarises the design parameters.

Table 6. Design parameters shared by the three configurations.

Element	Parameters
Backstepping	$(k_{p1}, k_{p2}, k_{p3}) = (7, 5, 2)$, $(k_{y1}, k_{y2}, k_{y3}) = (10, 10, 5)$, $\lvert U_i \rvert \leq 2.5$ V
Observer	$\bar{k} = [3, 3, 1]^T$, $\bar{\varepsilon} = 0.05$ rad, guard $\lvert D_i \rvert \geq 0.02$
Gain range	pitch $[4, 25]$, yaw $[3, 20]$
Rate limiter	$\rho_{\mathrm{up}} = 5$ s^{-1}, $\rho_{\mathrm{dn}} = 1$ s^{-1}, update period $T_s = 0.01$ s on rig and in simulation
Features	$q = 2\pi/2048$ rad, $n_q = 0.3$, $e_s = 0.008$ rad, $c_{\mathrm{ref}} = 0.1$ V

For each axis with tracking error $e(t)$, control voltage $U(t)$ and output-estimation error $\varepsilon(t)$, the following metrics are computed over the 100 s run with the true time vector: the root mean square error RMSE $= \left(\frac{1}{N}\sum e_k^2\right)^{1/2}$; the integral of time-weighted absolute error ITAE $= \int_0^T t|e|\,dt$; the control energy CE $= \int_0^T U^2 dt$; the total variation TV $= \sum_k |U_{k+1} - U_k|$, used as a chattering index; the percentage of samples at the voltage limit (SAT); the maximum absolute tracking error after the initial transient ($t \geq 20$ s); and the RMS value of ε for $t \geq 20$ s, used as a measure of estimation quality.

2.9.1. Statistical Evaluation

The interval $t \in [20, 95.4]$ s of each experiment contains six full periods (4π s) of the references. The metrics are computed on each period, reported as mean $\pm$ standard deviation over the six periods, and compared with Welch's t-test. These cycle-wise statistics quantify the variability within a run, not between runs, and are complemented by the offline replays and the Monte Carlo simulations described below.

2.9.2. Offline Replay

In a recording, the applied voltages U and the measured angles y fully determine the trajectory of the observer for any gain sequence, because the observer (8) is driven only by these two signals. Running the observer, the controllers and an alternative scheduler on the recorded (U, y) therefore gives exactly the estimates and the control commands that this scheduler would have produced on the same plant data. The feedback of the alternative commands on the plant is not reproduced, so that a replay compares estimation accuracy and command smoothness, not tracking. Replaying the logged scheduler reproduces the logged gains to within 10^{-14}, which validates the procedure. Each alternative is replayed on the three recordings, and the results are given as mean $\pm$ standard deviation over the recordings.

2.9.3. Simulation

The closed loop was simulated with the plant given by the manufacturer model [2], which differs from the design model (1) through signed thrust maps $a_i x|x| + b_i x$, a centrifugal term on the pitch axis, a lead-lag main-to-tail coupling $(3.5\text{ s} + 1)/(2\text{ s} + 1)$ and a main motor time constant of 1.2 s. The measured angles were quantised with the encoder step q and corrupted by white noise of 0.3 mrad standard deviation. Euler integration at 1 ms and a scheduler period of 10 ms were used, as on the rig. The simulation code implements the Simulink diagram; when driven by the recorded measurements, its observer, controller and scheduler routines reproduce the logged rig signals to within 10^{-13}, and in nominal conditions its metrics agree with those of the Simulink diagram to within 1.6%. Two scenarios were run with ten random seeds each. In scenario S1, the seeds draw a_1 uniformly within $\pm 10\%$ of its nominal value and the initial pitch angle within ± 0.05 rad. Scenario S2 adds a 40% reduction in the main-rotor thrust in reverse rotation, which reproduces

the asymmetry observed on the rig, and a 2 s disturbance at $t = 50$ s with accelerations of -1.5 rad/s^2 on pitch and 3 rad/s^2 on yaw. All schedulers see the same seeds, and differences are tested with paired t-tests.

3. Results

3.1. Simulation with Quantisation and Model Mismatch

In the absence of noise and quantisation, the features x_1 and x_3 vanish after the initial transient and each scheduler converges to a constant gain, so that a noise-free simulation cannot reveal the adaptation. The simulations are therefore run with quantised and noisy measurements. Figure 4 shows a representative run of scenario S2 and Figure 5 the corresponding gains. Tables 7 and 8 report the statistics over ten seeds, and Figure 6 summarises the relative changes with respect to IT3-CA.

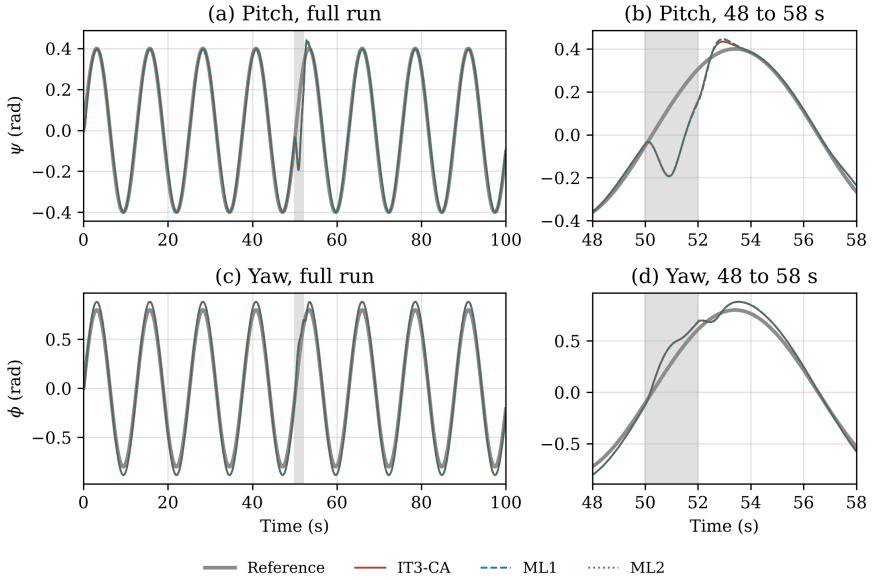

Figure 4. Simulated tracking in scenario S2 (seed 0): (**a**,**c**) full run; (**b**,**d**) enlarged view of the disturbance interval (shaded).

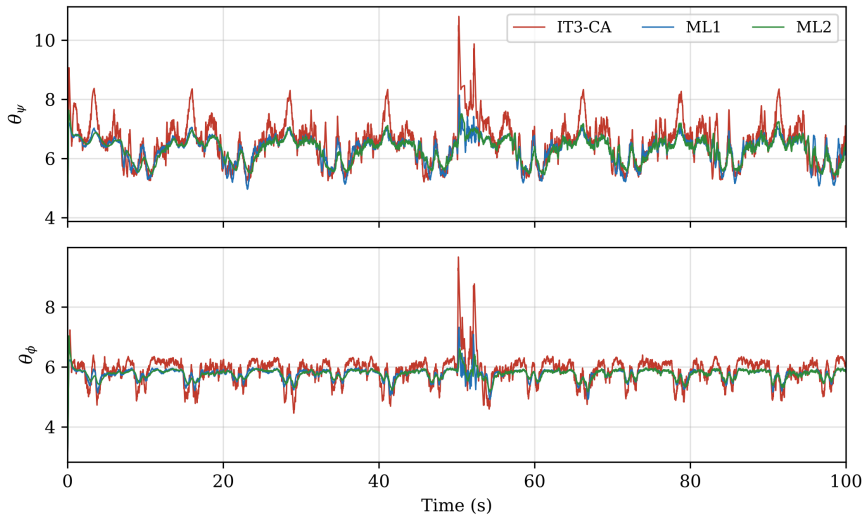

Figure 5. Scheduled observer gains in the simulation of Figure 4.

Table 7. Scenario S1, quantised measurements, nominal plant parameters with ±10% variation of a_1: mean ± standard deviation over ten seeds. IT2 is the scheduler of Remark 5 with $\mathcal{A} = \{1\}$; the fixed gain equals the mean IT3-CA gain of the same seed; RL is the rate-limited variant with $\rho_{\mathrm{up}} = \rho_{\mathrm{dn}} = 0.4\ \mathrm{s}^{-1}$. An asterisk marks a paired difference with IT3-CA at $p < 0.01$.

Metric	IT3-CA	ML1	ML2	IT2	IT3 Without x_3	Fixed Gain	RL
Pitch							
RMSE (mrad)	24.8 ± 0.6	25.0 ± 0.5	24.7 ± 0.5	24.8 ± 0.6	18.6 ± 0.6 *	22.9 ± 0.5 *	24.6 ± 0.4
ITAE (rad s^2)	82.2 ± 0.7	83.2 ± 0.7 *	82.4 ± 0.7	82.1 ± 0.5	62.1 ± 0.5 *	77.4 ± 0.7 *	82.3 ± 0.5
CE (V^2 s)	82.1 ± 0.8	81.7 ± 0.8	80.7 ± 0.9 *	81.9 ± 0.8	81.3 ± 1.7	80.3 ± 1.1 *	81.2 ± 1.2 *
TV (V)	825 ± 7	771 ± 7 *	770 ± 5 *	825 ± 5	1648 ± 13 *	849 ± 10 *	816 ± 7 *
RMS ε (mrad)	1.200 ± 0.009	1.214 ± 0.009 *	1.181 ± 0.007 *	1.198 ± 0.007	0.970 ± 0.005 *	1.088 ± 0.004 *	1.190 ± 0.004 *
Mean θ	6.52 ± 0.01	6.30 ± 0.01	6.31 ± 0.00	6.52 ± 0.02	8.85 ± 0.00	6.52 ± 0.01	6.51 ± 0.01
Yaw							
RMSE (mrad)	64.1 ± 0.1	64.1 ± 0.1 *	64.1 ± 0.1	64.1 ± 0.1	64.4 ± 0.1 *	64.1 ± 0.1 *	64.0 ± 0.1 *
ITAE (rad s^2)	290.4 ± 0.6	290.4 ± 0.6	290.3 ± 0.6	290.3 ± 0.5	292.2 ± 0.5 *	290.3 ± 0.6	290.3 ± 0.5
CE (V^2 s)	32.8 ± 0.2	32.6 ± 0.1 *	32.5 ± 0.1 *	32.8 ± 0.2	32.4 ± 0.2 *	32.3 ± 0.0 *	32.3 ± 0.1 *
TV (V)	761 ± 4	738 ± 3 *	737 ± 1 *	760 ± 5	1071 ± 5 *	756 ± 2 *	757 ± 3 *
RMS ε (mrad)	1.164 ± 0.004	1.158 ± 0.004	1.154 ± 0.003 *	1.164 ± 0.004	1.029 ± 0.002 *	1.123 ± 0.003 *	1.160 ± 0.002 *
Mean θ	5.85 ± 0.01	5.76 ± 0.01	5.76 ± 0.00	5.85 ± 0.01	6.93 ± 0.00	5.85 ± 0.01	5.85 ± 0.01

Table 8. Scenario S2, as S1 with a 40% weaker main-rotor thrust in reverse rotation and a 2 s disturbance at $t = 50$ s: mean ± standard deviation over ten seeds. An asterisk marks a paired difference with IT3-CA at $p < 0.01$.

Metric	IT3-CA	ML1	ML2	IT3 Without x_3	Fixed Gain	RL
Pitch						
RMSE (mrad)	35.6 ± 0.4	35.9 ± 0.4 *	35.9 ± 0.4 *	33.4 ± 0.4 *	35.7 ± 0.4	35.4 ± 0.4 *
ITAE (rad s^2)	74.8 ± 2.0	76.1 ± 2.0 *	76.0 ± 2.0 *	63.8 ± 1.3 *	74.6 ± 1.8 *	74.8 ± 1.9
CE (V^2 s)	128.3 ± 1.3	127.6 ± 1.2 *	127.8 ± 1.3 *	131.2 ± 1.6 *	128.3 ± 1.4	128.7 ± 1.4
TV (V)	797 ± 7	748 ± 7 *	750 ± 7 *	1614 ± 16 *	826 ± 4 *	790 ± 7 *
RMS ε (mrad)	1.109 ± 0.008	1.234 ± 0.008 *	1.246 ± 0.010 *	1.014 ± 0.004 *	1.300 ± 0.009 *	1.226 ± 0.023 *
Mean θ	6.51 ± 0.02	6.29 ± 0.01	6.29 ± 0.01	8.86 ± 0.00	6.51 ± 0.02	6.49 ± 0.02
Yaw						
RMSE (mrad)	61.3 ± 0.2	61.4 ± 0.2 *	61.4 ± 0.2 *	62.5 ± 0.2 *	61.6 ± 0.2 *	61.4 ± 0.2 *
ITAE (rad s^2)	275.0 ± 0.7	274.9 ± 0.8	274.9 ± 0.9	280.9 ± 0.7 *	275.8 ± 0.8 *	275.1 ± 0.8 *
CE (V^2 s)	43.1 ± 0.2	42.8 ± 0.2 *	42.9 ± 0.2 *	43.1 ± 0.3	42.7 ± 0.2 *	43.0 ± 0.1
TV (V)	765 ± 5	743 ± 3 *	743 ± 3 *	1068 ± 4 *	767 ± 5	760 ± 4 *
RMS ε (mrad)	1.369 ± 0.003	1.494 ± 0.007 *	1.533 ± 0.006 *	1.173 ± 0.003 *	1.539 ± 0.003 *	1.503 ± 0.011 *
Mean θ	5.87 ± 0.01	5.76 ± 0.00	5.76 ± 0.00	6.95 ± 0.00	5.87 ± 0.01	5.85 ± 0.01

Three observations follow. First, the chattering input is effective: without x_3 the scheduler operates at a gain about 36% higher in pitch and 18% higher in yaw, which improves the tracking and estimation errors but doubles the pitch total variation and increases the yaw total variation by about 40% in both scenarios. Second, reducing the IT3 sets to their IT2 core changes no metric significantly in S1 (all differences below 0.2%, $p \geq 0.08$), in agreement with Remark 5. Third, the effect of adaptation depends on the model mismatch. In S1, where the plant differs from the design model only through the structural terms listed above, the three schedulers give tracking errors within 1.2% of each other, and a fixed gain equal to the mean IT3-CA gain gives lower tracking and estimation errors at a 2.8% higher pitch total variation. In S2, IT3-CA gives an RMS estimation error 8 to 11% lower than the neural schedulers and 11 to 15% lower than the fixed mean gain,

with a slightly lower pitch tracking error, whereas the neural schedulers keep a 3 to 6% lower total variation. The rate-limited variant RL is equivalent to IT3-CA in S1, and in S2 it loses most of the estimation advantage of IT3-CA (RMS ε 10% higher), which quantifies the cost of satisfying the sufficient condition of Proposition 1.

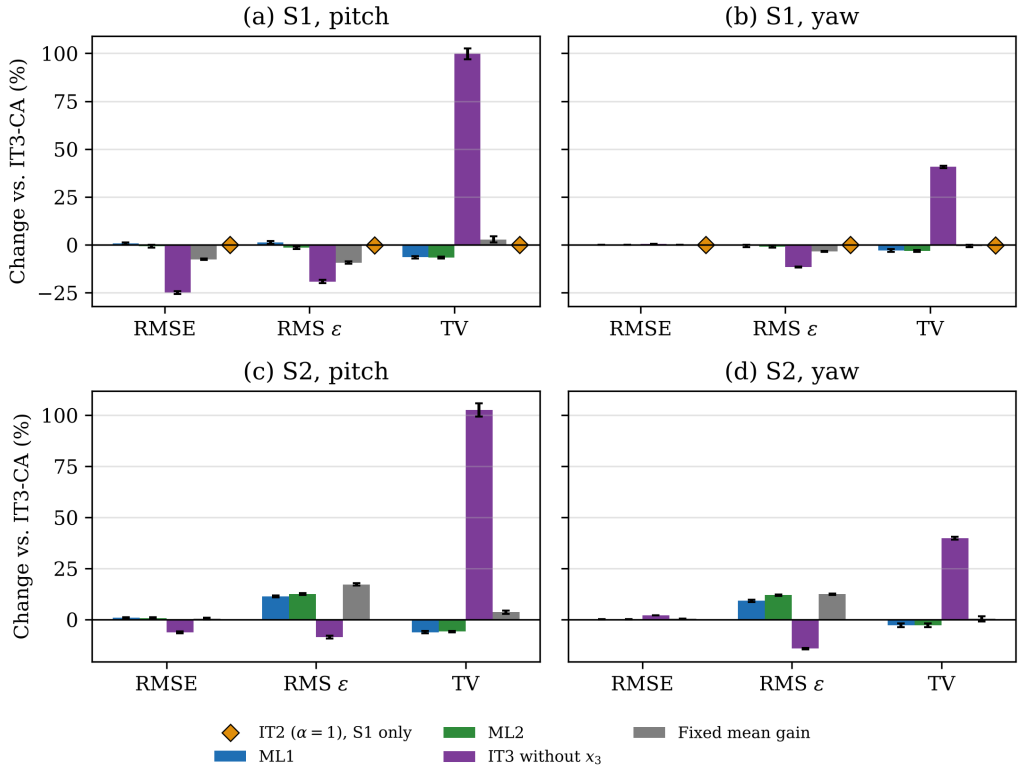

Figure 6. Change in the pitch and yaw metrics with respect to IT3-CA over ten seeds (mean ± standard deviation of the paired relative differences). (**a**,**b**) Scenario S1. (**c**,**d**) Scenario S2.

3.2. Real-Time Experiments

Figures 7 and 8 show the measured pitch and yaw responses of the three configurations on the physical TRMS, and Table 9 reports the metrics. For each configuration the pitch and yaw responses come from the same run, in which both axes tracked their references simultaneously; the performance of each axis therefore includes the effect of its coupling with the other. The three runs were recorded in the same session with the same update period of 10 ms for all schedulers. Table 10 reports the cycle-wise statistics.

Table 9. Real-time experimental results over the full runs (100 s, 1 kHz loop, schedulers every 10 ms).

Metric	Pitch			Yaw		
	IT3-CA	**ML1**	**ML2**	**IT3-CA**	**ML1**	**ML2**
RMSE (rad)	**0.0658**	0.1008	0.1012	0.0599	**0.0598**	0.0702
ITAE (rad s^2)	**248.4**	416.6	392.4	**233.9**	242.3	293.1
CE (V^2 s)	239.7	**222.4**	295.0	**242.9**	297.5	343.3
TV (V)	754	**640**	686	849	769	**738**
SAT (%)	23.5	**19.0**	31.2	**25.6**	32.6	39.9
Max $\lvert e \rvert$, $t \geq 20$ s (rad)	**0.172**	0.303	0.360	0.182	**0.150**	0.199
RMS ε (mrad)	**1.79**	2.28	2.54	**3.08**	4.49	4.56
Mean θ, $t \geq 20$ s	7.08	6.67	6.65	6.63	5.92	5.93

Bold values indicate the best result of each metric for each axis.

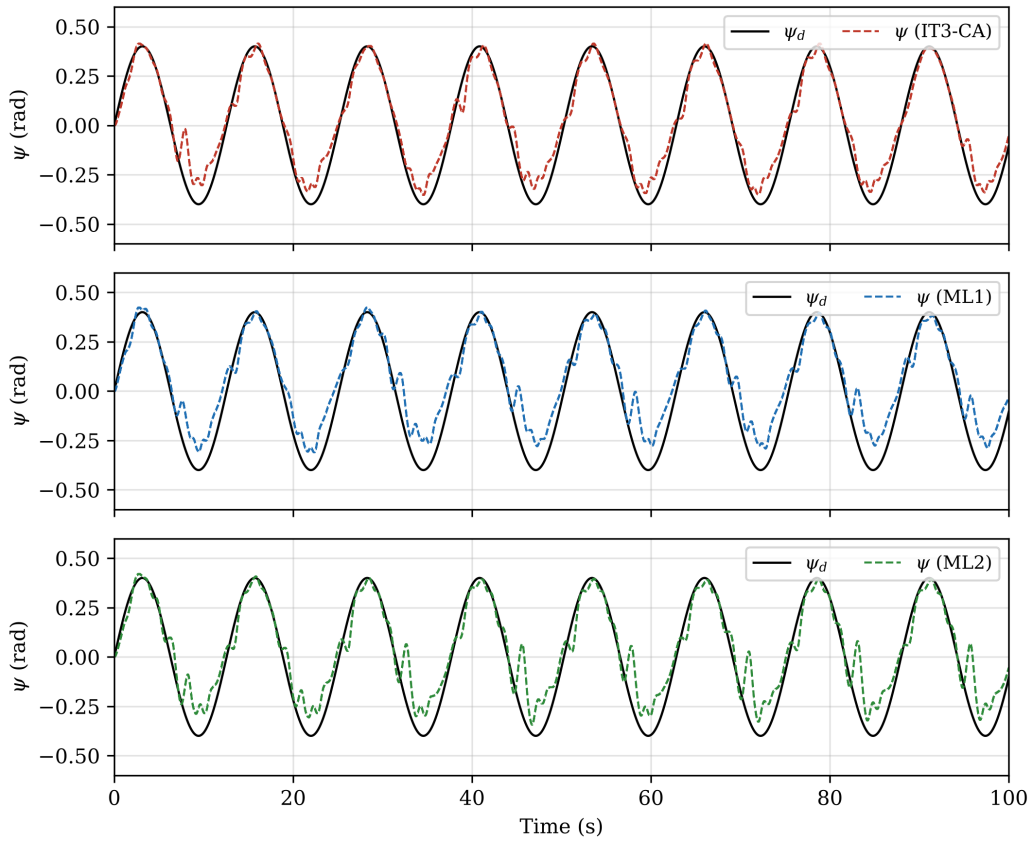

Figure 7. Real-time pitch tracking with the IT3-CA, ML1 and ML2 schedulers.

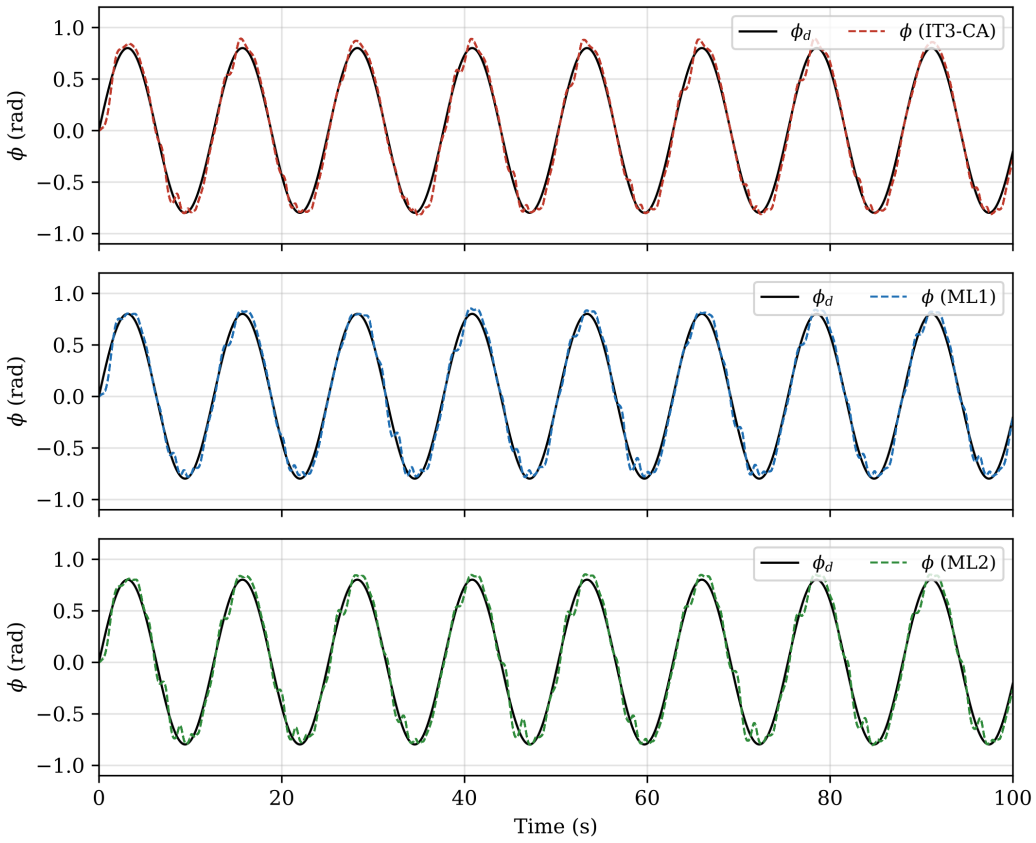

Figure 8. Real-time yaw tracking with the IT3-CA, ML1 and ML2 schedulers.

Table 10. Cycle-wise statistics of the real-time experiments: mean ± standard deviation over the six reference periods in $t \in [20, 95.4]$ s, and p-values of Welch's t-test against IT3-CA.

Metric per Cycle	IT3-CA	ML1	ML2	p (ML1)	p (ML2)
Pitch					
RMSE (mrad)	61.6 ± 3.9	101.4 ± 9.1	101.7 ± 11.7	2.8×10^{-5}	1.9×10^{-4}
RMS ε (mrad)	1.79 ± 0.10	2.33 ± 0.10	2.58 ± 0.12	2.4×10^{-6}	2.3×10^{-7}
TV (V)	93.9 ± 4.4	79.0 ± 3.5	85.5 ± 3.0	8.8×10^{-5}	0.004
CE (V^2 s)	30.3 ± 3.0	28.5 ± 1.5	37.6 ± 3.5	0.230	0.003
SAT (%)	23.4 ± 3.8	19.8 ± 3.3	31.3 ± 4.9	0.116	0.012
Yaw					
RMSE (mrad)	56.9 ± 5.8	58.0 ± 2.1	69.4 ± 6.5	0.666	0.006
RMS ε (mrad)	3.06 ± 0.10	4.48 ± 0.25	4.48 ± 0.12	8.0×10^{-6}	1.7×10^{-9}
TV (V)	107.1 ± 5.9	96.0 ± 2.8	92.3 ± 4.0	0.004	7.2×10^{-4}
CE (V^2 s)	29.4 ± 4.3	38.7 ± 1.6	43.3 ± 4.1	0.002	1.9×10^{-4}
SAT (%)	24.6 ± 5.7	34.2 ± 2.1	40.6 ± 5.5	0.007	5.9×10^{-4}

3.2.1. Tracking

The three configurations remain stable over the whole run. In pitch, IT3-CA gives the lowest RMSE and ITAE, with an RMSE 35% lower than both neural schedulers and a maximum steady-state error of 0.172 rad, against 0.303 and 0.360 rad. The cycle-wise difference is significant against both ($p < 0.001$). In yaw, IT3-CA and ML1 give the same RMSE within the cycle-to-cycle variability ($p = 0.67$), and both are 15% below ML2 ($p = 0.006$ for IT3-CA). In all three runs the largest pitch errors occur in the negative half-periods, where the measured pitch angle reaches about -0.30 to -0.35 rad instead of -0.4 rad while U_1 approaches its lower limit. This behaviour is common to the three schedulers and is discussed in Section 4.

3.2.2. Gain Adaptation and Estimation

On the rig, encoder quantisation and unmodelled dynamics keep the excess error and the chattering index active, and the schedulers do not converge to constants (Figure 9). IT3-CA operates at a mean gain of 7.08 in pitch and 6.63 in yaw, with excursions up to about 10.5 and 11 during transients, whereas the neural schedulers stay closer to their nominal gains. IT3-CA reduces the RMS output-estimation error to 1.79 mrad in pitch and 3.08 mrad in yaw, which is 22% and 30% lower in pitch, and 31% and 32% lower in yaw, than ML1 and ML2, respectively (Figure 10). These differences are significant over the six cycles ($p < 10^{-5}$).

The mechanism behind the IT3-CA gain is visible in Figure 11, which shows the three features recomputed offline from the logged estimation error and control voltage; feeding these logged signals to the scheduler reproduces the logged gain to within 10^{-14}. Each burst of the excess error x_1 with a positive trend x_2 produces a gain peak. With the scheduler executed at its design period, the chattering index x_3 is now active: it exceeds 0.25 during 84% of the time in pitch and 78% in yaw (mean values 0.51 and 0.41), so that the High-chattering branch of the rule base lowers the gain whenever the control voltage starts to chatter, as seen for instance around 46 to 49 s and 58 to 60 s in pitch.

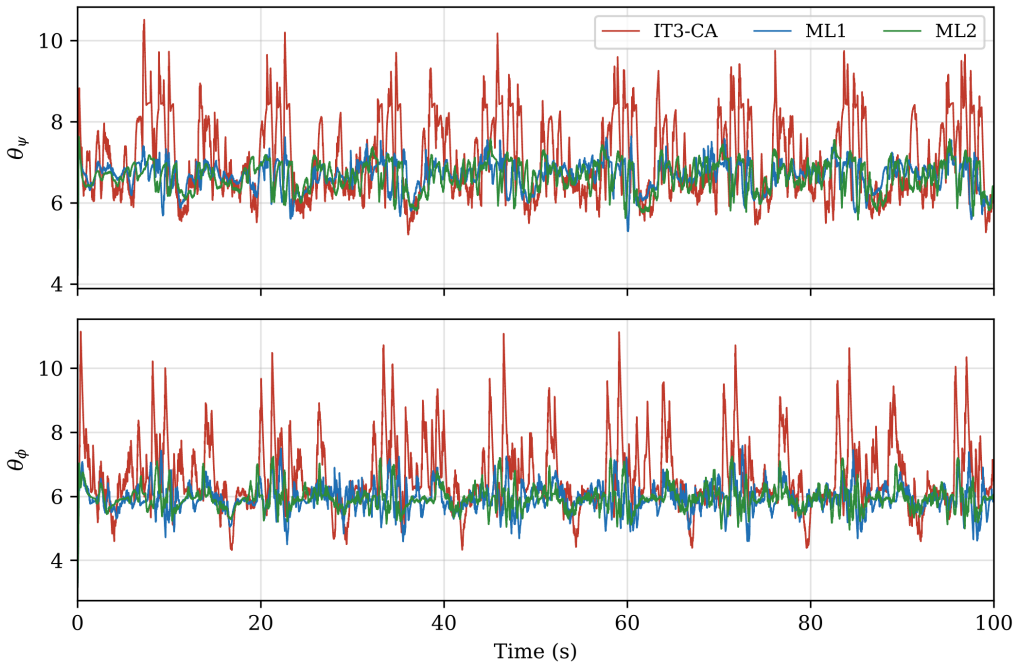

Figure 9. Scheduled observer gains during the real-time experiments.

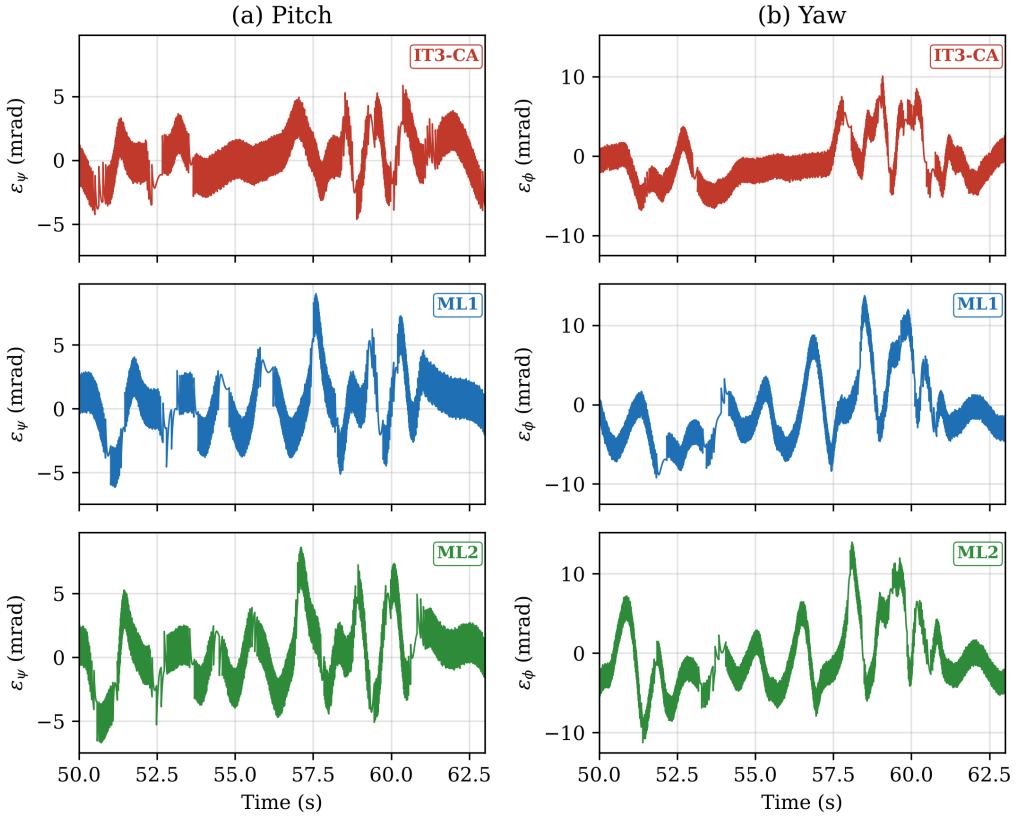

Figure 10. Output-estimation errors during the real-time experiments, detail between 50 and 63 s, one panel per scheduler. (**a**) Pitch. (**b**) Yaw.

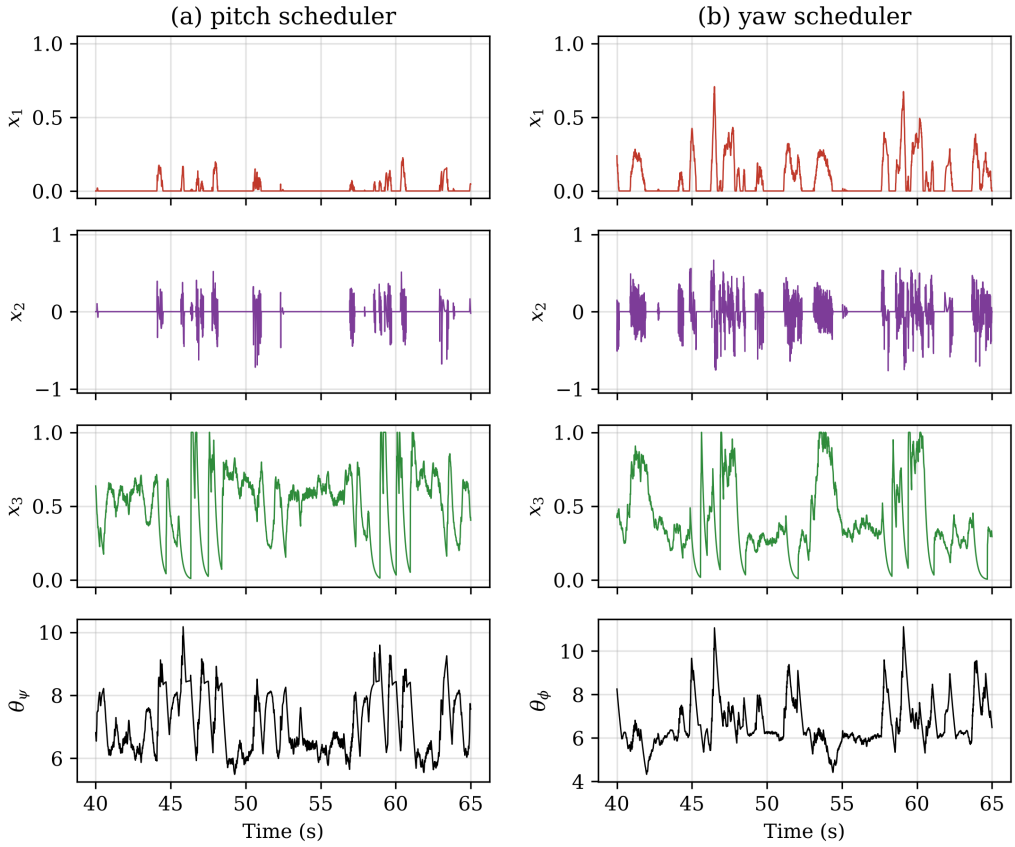

Figure 11. IT3-CA features and scheduled gain during the real-time experiment, detail between 40 and 65 s: excess estimation error x_1, its trend x_2, chattering index x_3 and observer gain. (**a**) Pitch. (**b**) Yaw.

3.2.3. Control Effort and Chattering

The price of the higher gain appears in the control signals (Figure 12). The total variation in IT3-CA is 18% and 10% above that of ML1 and ML2 in pitch, and 10% and 15% above in yaw. The cycle-wise differences are significant, but the ratio is much smaller than when the fuzzy scheduler was executed at the control rate, where its chattering input was inactive. The control energy and the saturation ratio of IT3-CA are the lowest in yaw ($p \leq 0.007$), and in pitch they lie between those of ML1 and ML2. Figure 13 summarises the trade-off by normalising every cycle-wise metric to the IT3-CA value.

3.3. Ablations on the Recorded Data

The offline replay isolates the contribution of each element of the scheduler on the plant data actually recorded. Table 11 reports, for each alternative, the RMS estimation error, the total variation in the control command and the mean gain over the three recordings, and Figure 14 places them in the plane of estimation accuracy and command smoothness, together with the trade-off curves obtained by sweeping $c_{\text{ref}} \in \{0.05, 0.1, 0.2, 0.4\}$ V for IT3-CA and $w_c \in \{0, 0.125, 0.5, 2\}$ for the neural schedulers.

3.3.1. Third Fuzzy Dimension

The IT2 variant reproduces IT3-CA to within 0.1% on every recording and every metric. On this scheduler, the third dimension of the sets therefore brings no measurable benefit, as anticipated in Remark 5.

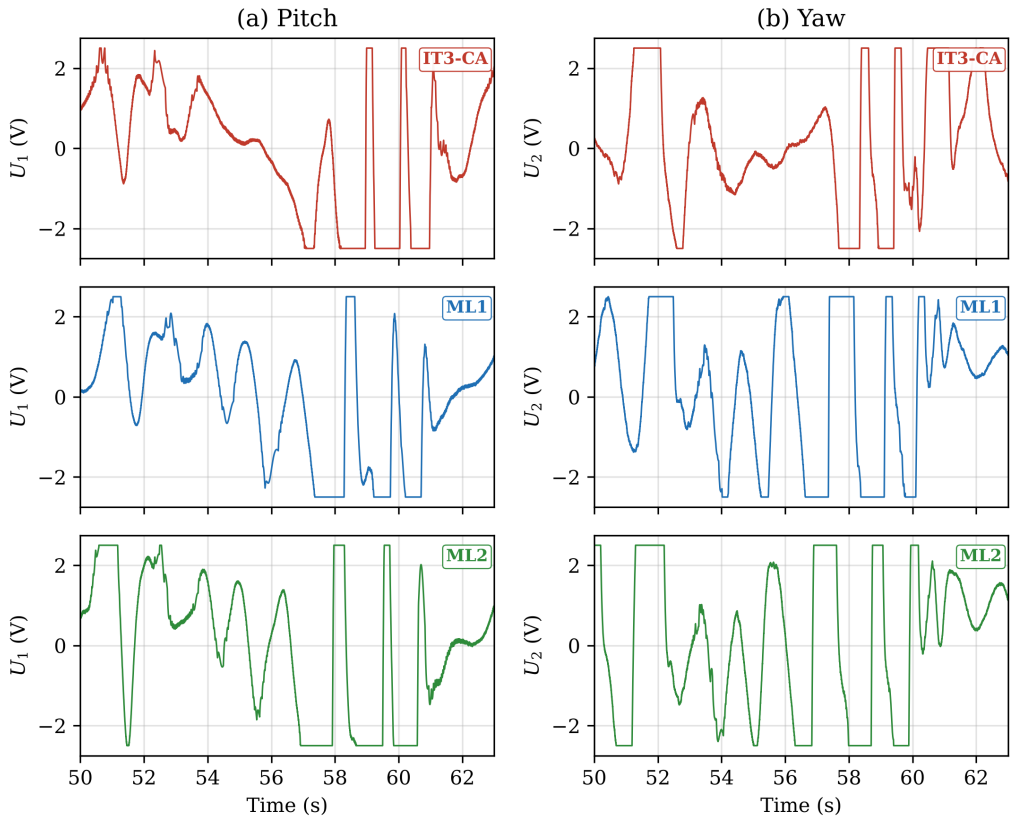

Figure 12. Control voltages during the real-time experiments, detail between 50 and 63 s, one panel per scheduler. (**a**) Pitch. (**b**) Yaw.

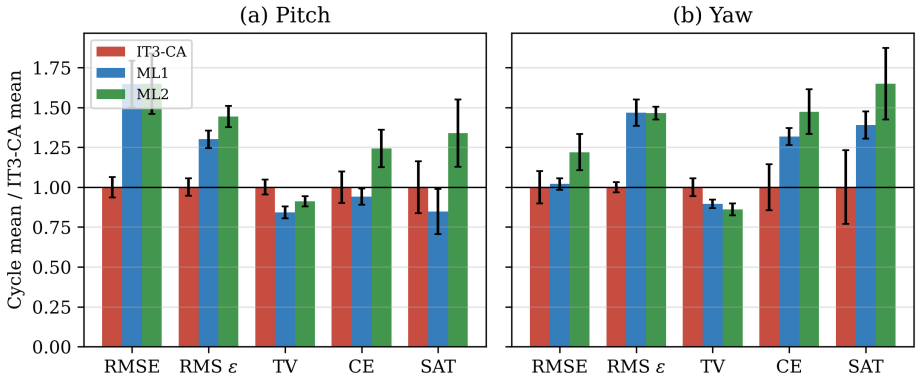

Figure 13. Cycle-wise experimental metrics normalised to the IT3-CA mean (bars: mean over six cycles; whiskers: standard deviation). Values below one favour the neural scheduler. (**a**) Pitch. (**b**) Yaw.

3.3.2. Chattering Input

Removing x_3 raises the mean gain by about 1.5 in pitch and 0.7 in yaw and lowers the estimation error by 21% and 15% on average, but increases the total variation in the command by 26 to 45% in pitch and by 5 to 16% in yaw across the recordings. The chattering input is thus the element that keeps the scheduled gain from trading smoothness for accuracy without limit.

Table 11. Offline replay of the three recordings ($t \geq 20$ s): mean $\pm$ standard deviation over the recordings, and mean relative change with respect to IT3-CA. The fixed gain is set, for each recording, to the mean IT3-CA gain of the same replay.

Scheduler	RMS ε (Mrad)		TV of U (V)		Mean θ	
	Pitch	Yaw	Pitch	Yaw	Pitch	Yaw
IT3-CA	1.83 ± 0.09	3.02 ± 0.08	553 ± 88	634 ± 47	7.43	6.80
IT2 ($\mathcal{A} = \{1\}$)	1.83 ± 0.09 (0.0%)	3.02 ± 0.08 (0.0%)	553 ± 88 (0.0%)	634 ± 47 (0.0%)	7.43	6.80
IT3 without x_3	1.45 ± 0.09 (-21%)	2.56 ± 0.05 (-15%)	742 ± 147 (+34%)	695 ± 87 (+10%)	8.91	7.53
Fixed mean gain	2.00 ± 0.17 (+9%)	3.78 ± 0.06 (+25%)	582 ± 109 (+5%)	694 ± 51 (+9%)	7.43	6.80
RL, $\rho = 0.4$ s^{-1}	2.02 ± 0.17 (+10%)	3.95 ± 0.06 (+31%)	554 ± 89 (+0.1%)	633 ± 36 (-0.2%)	7.32	6.45
ML1	2.34 ± 0.22 (+28%)	4.49 ± 0.06 (+49%)	482 ± 118 (-13%)	608 ± 18 (-4%)	6.63	5.90
ML2	2.34 ± 0.22 (+28%)	4.49 ± 0.06 (+49%)	481 ± 115 (-13%)	612 ± 17 (-3%)	6.63	5.91

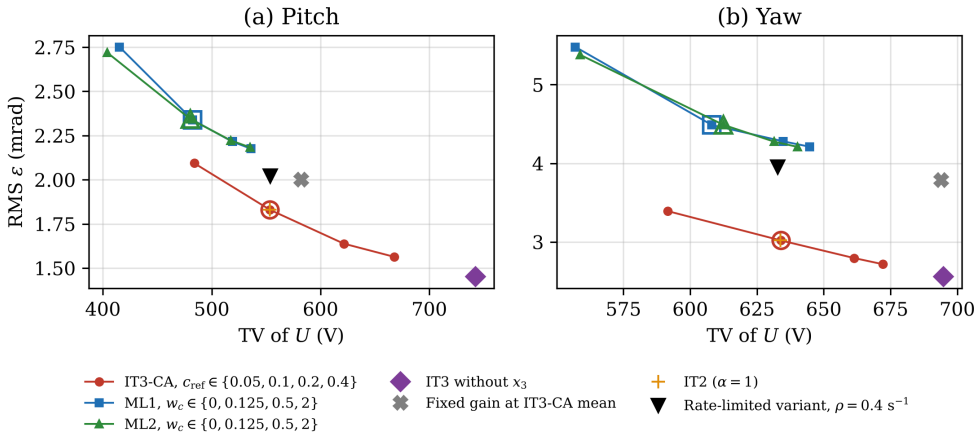

Figure 14. Estimation accuracy versus command smoothness on the replayed recordings (mean over three recordings). Lines: trade-off curves of IT3-CA (c_{ref} sweep) and of the neural schedulers (w_c sweep); large open markers: nominal settings. (**a**) Pitch. (**b**) Yaw.

3.3.3. Adaptation Versus Mean Gain

A constant gain equal to the mean IT3-CA gain is worse than IT3-CA in both estimation and smoothness on every recording: the estimation error is 6 to 14% higher in pitch and 24 to 27% higher in yaw, and the total variation 1 to 10% higher. The benefit of IT3-CA is therefore not a mean-gain effect: raising the gain during error bursts and lowering it when the command chatters outperforms the constant gain that has the same average. The rate-limited variant, which satisfies the sufficient condition of Proposition 1, keeps the smoothness of IT3-CA but loses about one third of its estimation advantage over the neural schedulers in pitch and about two thirds in yaw.

3.3.4. Tuning

Over the swept ranges, the IT3-CA curve lies below the curves of both neural schedulers in both axes (Figure 14). At equal total variation, for instance with $c_{\mathrm{ref}} = 0.05$ V in pitch, IT3-CA gives an estimation error about 11% lower than the neural schedulers at their nominal setting, and the neural schedulers do not reach the estimation error of IT3-CA even without the chattering penalty ($w_c = 0$), because the nominal-gain term of the cost (19) keeps their gain close to θ_0. The ranking of the families is therefore not an artefact of a particular tuning point within the explored ranges. The two neural schedulers behave almost identically on the replays, which indicates that their difference on the rig stems mainly from their transient learning rather than from the network structure.

3.4. Execution Time

The three schedulers were implemented in C, compiled with optimisation level O2 and timed on one core of a 2.1 GHz Intel Xeon processor. The C implementations reproduce the scheduler blocks to within 10^{-14} on the recorded features. Each scheduler was fed with the recorded pitch error and voltage, and 2×10^5 updates were timed individually. Table 12 reports the results. The 99th percentile of the IT3-CA update is 2.6 µs, which represents 0.26% of the 1 ms control step and 0.03% of the 10 ms scheduler period; for both axes the load at 100 Hz is below 0.05%. The rare maxima, observed for all three schedulers including ML2, whose arithmetic is fixed, are caused by pre-emption of the general-purpose host operating system and not by the algorithms, whose loops are bounded. The additional cost of IT3-CA, about four times that of ML1 and seven times that of ML2, is therefore negligible for the real-time target.

Table 12. Execution time per scheduler update (one axis), compiled C implementation, 2×10^5 timed updates.

Scheduler	Mean (µs)	Median (µs)	99th Perc. (µs)	99.99th Perc. (µs)	Load at 100 Hz (%)
IT3-CA	1.70	1.54	2.61	57.5	0.017
ML1 (MLP-ADAM)	0.41	0.35	0.57	81.2	0.004
ML2 (RBF-NLMS)	0.25	0.23	0.40	16.6	0.003

3.5. Comparison with the Previous IT2 Fuzzy Scheme

The same TRMS and interface were used in our previous study with an IT2 Takagi–Sugeno fuzzy supervisor of the observer gain [20], which reported RMSE values of 0.0883 rad (pitch) and 0.1979 rad (yaw) and ITAE values of 344.28 and 834.90. The complete scheme proposed here gives RMSE values 25% and 70% lower and ITAE values 28% and 72% lower. This comparison concerns complete schemes: the observer, the controller tuning, the loop rate and the scheduler differ between the two studies, so that the improvement cannot be attributed to the fuzzy scheduler alone, and the large yaw improvement is shared by the two neural schedulers (Table 9), which points to the observer correction and the yaw controller tuning. The contribution of the fuzzy architecture itself is isolated by the ablation of Section 3.3, in which the observer, the controller, the inputs, the gain limits and the update period are identical and only the fuzzy sets change: IT3 and IT2 sets give the same results within 0.1%.

4. Discussion

The results reveal a consistent trade-off between estimation accuracy and control smoothness, and they identify which elements of the scheduler move along it. The fuzzy scheduler observes the excess estimation error, its trend and the chattering of the command, raises the gain when the error rises above the measured noise floor, and lowers it when the command chatters. Executed at its design period, it achieved the most accurate observer on hardware and the best pitch tracking, with a total variation 10 to 18% above that of the neural schedulers. The neural schedulers descend a cost in which the distance to a nominal gain plays a stabilising role; they remain close to that gain, which yields slightly smoother commands at the expense of a larger estimation error. The offline replays show that this ranking holds over the explored tuning ranges of both families.

The ablations clarify the source of the benefit. The chattering input is essential: without it the scheduler drifts to a higher gain, and the total variation increases by up to 45% on the replays and up to 100% in simulation. Adaptation itself matters under model mismatch: a constant gain with the same average is worse in both estimation and smoothness on the

rig recordings and in the mismatched simulation scenario, although it is sufficient when the plant is close to the design model. The third dimension of the fuzzy sets, by contrast, has no measurable effect in this configuration, because its symmetric spread cancels to first order in the type-reduced output. The relevant novelty of the scheduler is therefore its scheduling information, namely the noise-floor normalisation and the chattering index, rather than the order of its fuzzy sets. An asymmetric spread, an uncertainty that depends on the operating point, or a learned spread could give the third dimension a role, and this question remains open.

The stability analysis has been sharpened but remains sufficient. With an optimised Lyapunov matrix and an averaged-rate formulation, Proposition 1 certifies, on the recorded operating envelope, any scheduler whose relative gain rate stays below 0.45 s^{-1}, and it explains why: the decay of the scaled error, which grows with the gain, has to dominate the variation in the scaling itself. The recorded schedulers change the gain faster, so that their stability on the rig is an experimental observation, and a rate-limited scheduler that satisfies the condition loses part of the estimation benefit. A less conservative certificate, for instance with a gain-dependent Lyapunov matrix or a dwell-time argument exploiting the burst structure of the gain, would narrow this gap.

The experiments share a pitch error concentrated in the negative half-periods. The main rotor of the TRMS produces a markedly weaker thrust in reverse rotation than the quadratic map F_1 assumes, and the controller saturates the main motor when it tries to drive the beam below the counterweight equilibrium. The resulting error is independent of the scheduler, as confirmed by scenario S2, and would be reduced by an asymmetric thrust model in the controller and the observer. Such a signed map would also make D_1 bounded below by b_1, so that the guard of Remark 1 would no longer be needed.

The study has several limitations. First, the experimental comparison rests on one 100 s run per scheduler; the cycle-wise statistics quantify the variability within a run, and between-run variability is addressed indirectly by the replays over three recordings and by the Monte Carlo simulations, but repeated hardware runs with randomised order would strengthen the conclusions. Second, the schedulers were tuned by design reasoning and compared through one-parameter sweeps, not through a joint optimisation under an identical budget. Third, the replays reproduce the observer and the commands on the recorded plant data but not the closed-loop effect of the alternative commands. Finally, the execution times were measured on a general-purpose processor, not on the real-time target itself.

5. Conclusions

A chattering-aware fuzzy scheduler of the high-gain observer gain was proposed for the output-feedback backstepping control of a Twin Rotor MIMO System, together with a corrected observer with saturated injection and a practical-stability condition for the time-varying gain that is evaluated along the recorded operating points. Two neural schedulers sharing the same features and gain mapping, an ADAM-trained MLP and an NLMS-trained RBF network, were designed as benchmarks. In real-time experiments with simultaneous pitch and yaw tracking and a common scheduler period, the fuzzy scheduler gave a pitch tracking RMSE 35% lower and RMS estimation errors 22 to 32% lower than both neural schedulers, with statistically significant differences over six reference cycles, at the cost of a 10 to 18% higher control total variation. Ablations on the recorded data and quantised Monte Carlo simulations showed that the chattering input limits the total variation by up to a factor of two, that the adaptation outperforms a constant gain with the same mean under model mismatch, and that the interval type-3 sets behave as their interval type-2 core. In a compiled implementation, 99% of the scheduler updates take less than 3 µs.

https://doi.org/10.3390/act15120621

Future work will address repeated hardware trials with randomised order, a joint optimisation of the fuzzy and neural schedulers under an identical budget, a gain-dependent Lyapunov analysis that certifies faster gain variations, scheduling configurations in which the third fuzzy dimension carries operating-point-dependent uncertainty, and an asymmetric thrust model of the main rotor.

Author Contributions: Conceptualization, A.S., A.B. and A.K.; methodology, A.S., A.B. and A.K.; software, A.B. and A.K.; validation, A.S., A.B., A.K. and A.A.; formal analysis, A.S., A.K. and M.H.Z.; investigation, A.S., A.B. and M.H.Z.; resources, A.S., A.A. and J.-J.T.; data curation, A.S., A.B. and A.A.; writing, original draft preparation, A.S., A.B. and A.K.; writing, review and editing, A.S., A.B., A.K., A.A., M.H.Z. and J.-J.T.; visualisation, A.B. and A.K.; supervision, A.S. and A.K.; project administration, A.S., 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: This research received no external funding. The APC was funded by Multimedia University, Malaysia.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The experimental and simulation data, the controller, observer and scheduler source files, the replay and timing programs, 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. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

ADAM	Adaptive moment estimation
CE	Control energy
HGO	High-gain observer
IT2	Interval type-2
IT3	Interval type-3
IT3-CA	Chattering-aware interval type-3 scheduler
ITAE	Integral of time-weighted absolute error
KM	Karnik–Mendel
LMI	Linear matrix inequality
MIMO	Multiple-input multiple-output
MLP	Multilayer perceptron
NLMS	Normalised least mean squares
RBF	Radial basis function
RL	Rate-limited variant of IT3-CA
RMSE	Root mean square error
SAT	Saturation ratio
TRMS	Twin Rotor MIMO System
TV	Total variation

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.
2. Feedback Instruments Ltd. *Twin Rotor MIMO System Control Experiments 33-949S*; Feedback Instruments Ltd.: Crowborough, UK, 2006.

3. Irfan, S.; Zhao, L.; Ullah, S.; Javaid, U.; Iqbal, J. Differentiator- and Observer-Based Feedback Linearized Advanced Nonlinear Control Strategies for an Unmanned Aerial Vehicle System. *Drones* **2024**, *8*, 527.
4. Irfan, S.; Zhao, L.; Javaid, U.; Iqbal, J. URED Observer-Based Feedback Linearized Neuro Adaptive SMC for a Twin Rotor MIMO System: Design and Experimental Study. *Guid. Navig. Control* **2025**, *5*, 559–583. https://doi.org/10.1142/S2737480725500281.
5. 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.
6. 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.
7. Beloufa, A.; Kacimi, A.; Tahraoui, S.; Senoussaoui, A.; Azzouz, A.; Zaid, M.H.; Tiang, J.-J. 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. *Actuators* **2026**, *15*, 453.
8. Krstic, M.; Kanellakopoulos, I.; Kokotovic, P.V. *Nonlinear and Adaptive Control Design*; Wiley: New York, NY, USA, 1995.
9. Esfandiari, F.; Khalil, H.K. Output feedback stabilization of fully linearizable systems. *Int. J. Control* **1992**, *56*, 1007–1037.
10. Khalil, H.K. *Nonlinear Systems*, 3rd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 2002.
11. Khalil, H.K.; Praly, L. High-gain observers in nonlinear feedback control. *Int. J. Robust Nonlinear Control* **2014**, *24*, 993–1015.
12. Gauthier, J.P.; Hammouri, H.; Othman, S. A simple observer for nonlinear systems: Applications to bioreactors. *IEEE Trans. Autom. Control* **1992**, *37*, 875–880.
13. 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.
14. Vasiljevic, L.K.; Khalil, H.K. Error bounds in differentiation of noisy signals by high-gain observers. *Syst. Control Lett.* **2008**, *57*, 856–862.
15. Ahrens, J.H.; Khalil, H.K. High-gain observers in the presence of measurement noise: A switched-gain approach. *Automatica* **2009**, *45*, 936–943.
16. Boizot, N.; Busvelle, E.; Gauthier, J.P. An adaptive high-gain observer for nonlinear systems. *Automatica* **2010**, *46*, 1483–1488.
17. Sanfelice, R.G.; Praly, L. On the performance of high-gain observers with gain adaptation under measurement noise. *Automatica* **2011**, *47*, 2165–2176.
18. Prasov, A.A.; Khalil, H.K. A nonlinear high-gain observer for systems with measurement noise in a feedback control framework. *IEEE Trans. Autom. Control* **2013**, *58*, 569–580.
19. Astolfi, D.; Marconi, L. A high-gain nonlinear observer with limited gain power. *IEEE Trans. Autom. Control* **2015**, *60*, 3059–3064.
20. 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.
21. 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.
22. Mohammadzadeh, A.; Sabzalian, M.H.; Zhang, W. An interval type-3 fuzzy system and a new online fractional-order learning algorithm: Theory and practice. *IEEE Trans. Fuzzy Syst.* **2020**, *28*, 1940–1950.
23. Castillo, O.; Castro, J.R.; Melin, P. *Interval Type-3 Fuzzy Systems: Theory and Design*; Springer: Cham, Switzerland, 2022.
24. Liu, Z.; Mohammadzadeh, A.; Turabieh, H.; Mafarja, M.; Band, S.S.; Mosavi, A. A new online learned interval type-3 fuzzy control system for solar energy management systems. *IEEE Access* **2021**, *9*, 10498–10508.
25. Qasem, S.N.; Ahmadian, A.; Mohammadzadeh, A.; Rathinasamy, S.; Pahlevanzadeh, B. A type-3 logic fuzzy system: Optimized by a correntropy based Kalman filter with adaptive fuzzy kernel size. *Inf. Sci.* **2021**, *572*, 424–443.
26. Mohammadzadeh, A.; Taghavifar, H.; Zhang, C.; Alattas, K.A.; Liu, J.; Vu, M.T. A non-linear fractional-order type-3 fuzzy control for enhanced path-tracking performance of autonomous cars. *IET Control Theory Appl.* **2024**, *18*, 40–54.
27. Mendel, J.M.; Liu, F.; Zhai, D. α-Plane representation for type-2 fuzzy sets: Theory and applications. *IEEE Trans. Fuzzy Syst.* **2009**, *17*, 1189–1207.
28. Wagner, C.; Hagras, H. Toward general type-2 fuzzy logic systems based on zSlices. *IEEE Trans. Fuzzy Syst.* **2010**, *18*, 637–660.
29. Karnik, N.N.; Mendel, J.M. Centroid of a type-2 fuzzy set. *Inf. Sci.* **2001**, *132*, 195–220.
30. Mendel, J.M. *Uncertain Rule-Based Fuzzy Systems: Introduction and New Directions*, 2nd ed.; Springer: Cham, Switzerland, 2017.
31. Vindigni, C.R.; Esposito, A.; Orlando, C.; Alaimo, A. Comparison of Piezoelectric Stack-Based Passive and Active Vibration Suppression Systems for Satellite Solar Panels. *Vibration* **2025**, *8*, 15.
32. Pan, T.; Shen, X. A novel active and passive integrated vibration isolator based on piezoelectric stack: Characteristic analysis and experimental investigation. *J. Vib. Control* **2025**, *31*, 592–607.
33. Hou, Q.; Wang, H.; Lee, C.H.T.; Ding, S. Composite Adaptive Super-Twisting Sliding Mode Control Using Barrier Function for PM Motor Drives Toward Electric Aircraft Applications. *IEEE Trans. Power Electron.* **2025**, *40*, 16255–16264.

34. Hou, Q.; Wang, H.; Zhao, C.; Xu, S.; Zuo, Y.; Lee, C.H.T.; Ding, S. Super-Twisting Extended State Observer-Based Quasi-Proportional-Resonant Controller for Permanent Magnet Synchronous Motor Drive System. *IEEE Trans. Transp. Electrif.* **2024**, *10*, 1596–1604.
35. Jiang, L.; Ding, S.; Hou, Q.; Ding, C.; Chen, X. Global Fixed-Time SOSM Constant-Voltage Control for IPT–Buck Heating Systems Under Coil Misalignment and Load Variations. *IEEE Trans. Circuits Syst. II Express Briefs* **2026**, *early access*. https://doi.org/10.1109/TCSII.2026.3745863.
36. Boyd, S.; El Ghaoui, L.; Feron, E.; Balakrishnan, V. *Linear Matrix Inequalities in System and Control Theory*; SIAM: Philadelphia, PA, USA, 1994.
37. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations, San Diego, CA, USA, 7–9 May 2015.
38. Park, J.; Sandberg, I.W. Universal approximation using radial-basis-function networks. *Neural Comput.* **1991**, *3*, 246–257.
39. Haykin, S. *Adaptive Filter Theory*, 5th ed.; Pearson: Upper Saddle River, NJ, USA, 2014.

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.