electronics

Article

ZMP-Guided Ground Anchoring for Quadrotor Landings Using MRAC-Neural Networks

Özgür Altınışık [1,*] and Seta Bogosyan [2]

[1] Department of Mechatronics Engineering, Istanbul Technical University, 34469 Istanbul, Turkey
[2] Department of Electrical and Electronics Engineering, MEF University, 34396 Istanbul, Turkey; estradao@mef.edu.tr
* Correspondence: altinisik19@itu.edu.tr

Abstract

Landing underactuated quadrotors on inclines is challenging as sudden contact transients convert translational kinetic energy into critical edge-tipping moments. Conventional controllers struggle to mitigate these sub-second impacts without computationally intensive prior terrain models or payload-restricting mechanical shock absorbers. This paper introduces a stabilization approach combining a Zero Moment Point (ZMP)-guided anchoring strategy with an impact-resilient Model Reference Adaptive Control Neural Network (MRAC-NN). The ZMP serves as a proactive geometric threshold. Evaluating overturning moments directly in the body frame before physical tilt occurs triggers immediate asymmetric bidirectional thrust. This artificially augments the surface normal force, thereby securing the vehicle on steep inclines that exceed natural static friction limits. Concurrently, the computationally efficient MRAC-NN suppresses unstructured aerodynamic disturbances and high-frequency impact shocks. To prevent parameter wind-up during sudden kinematic arrest at touchdown, the neuro-adaptive law is structurally augmented with robust bounding and dynamic state-locking mechanisms, mathematically guaranteeing Uniform Ultimate Boundedness (UUB). The framework is validated through high-fidelity simulations that incorporate Linear Complementarity Problem (LCP) rigid-impact constraints and stick–slip friction. Results demonstrate the architecture effectively suppresses post-impact bouncing and transforms critical tipping moments into controlled planar slides, reducing stabilization penalties by up to 47.5% and expanding the survivable flight envelope to 60° inclines where classical methods fail.

Keywords: quadrotor impact stabilization; Zero Moment Point (ZMP); ground anchoring; Model Reference Adaptive Control (MRAC); adaptive neural networks; Linear Complementarity Problem (LCP)

Academic Editors: Mohammad Jafari and Mahmut Reyhanoglu

Received: 22 July 2026
Revised: 5 September 2026
Accepted: 6 September 2026
Published: 11 September 2026

1. Introduction

Landing a quadrotor on an inclined surface poses a critical risk of structural rollover. Unlike steady-state hovering, the touchdown sequence forces a rapid transition from a continuous free-flight manifold to a strictly constrained geometric state. Because landing gears rarely strike the surface simultaneously, vertical kinetic energy is abruptly converted into extreme multi-axis tipping moments. Mastering these aggressive contact dynamics is crucial not only for individual multirotor stabilization but also for broadening the operational scope of autonomous vehicles. As aerial platforms increasingly transition toward collaborative swarm applications [1] and distributed model predictive control

 https://doi.org/10.3390/electronics15184130

across nano-UAV networks [2], securing individual vehicle survivability against unmodeled contact shocks in complex environments becomes a foundational prerequisite.

Addressing these sub-second impact dynamics exposes the structural limitations of conventional continuous control paradigms. Classical linear controllers (e.g., LQR, PID) [3] structurally fail to adapt to sudden center-of-mass shifts induced by asymmetric single-leg impacts due to their constant-inertia assumptions. Although Model Predictive Control (MPC) [4] theoretically excels at constraint handling, computing non-differentiable rigid-impact forces in real time overburdens standard flight controllers. Its application relies on computationally intensive prior surface models; thus, without perfectly accurate terrain maps, strict predictive horizons frequently destabilize. Hardware augmentations like adaptive landing gears [5] absorb kinetic shocks but impose predictable-incline constraints and payload-restricting weight penalties. Conversely, offline data-driven approaches such as the "Neural Lander" [6] mitigate complex aerodynamic effects but fundamentally lack real-time adaptation to sudden payload shifts, sensor degradation, or out-of-distribution environmental disturbances.

To overcome these limitations under closed-loop conditions, recent literature has increasingly focused on hybrid architectures that combine analytical model-based control with neural learning. For instance, integrating robust model-based strategies, such as Linear Quadratic Control (LQC), with asymmetric neural temporal models—specifically conditional LSTM autoencoders (CLSTM-AE)—has proven highly effective at characterizing closed-loop behavior and providing robust adaptive decision mechanisms against measurement noise and system variability [7]. While such temporal diagnostic frameworks are primarily designed for fault detection and isolation, where feedback compensation might otherwise mask fault signatures, they validate the broader need to embed neural learning mechanisms within nonlinear control loops to handle severe, unmodeled state deviations.

To achieve generalized stabilization without specialized landing hardware, prior surface knowledge, or computationally intensive predictive solvers, the controller must actively react to the physical boundaries of contact. Because standard cascaded flight controllers treat ground impacts as generic disturbances, translating an edge-tipping displacement into a corrective attitude reference introduces significant phase lag, often causing rollover before actuation fully takes effect. The Zero Moment Point (ZMP), originally formalized for bipedal locomotion [8,9], directly addresses this tipping problem by bypassing this lag. Although traditionally applied to ground robotics, an asymmetrically grounded quadrotor momentarily behaves as an underactuated inverted pendulum, pivoting about its initial contact point, thereby making bipedal stability criteria mathematically applicable. Mapping the exact coordinate where horizontal overturning moments vanish allows the controller to evaluate the imminent rollover threat directly from the ground reaction forces before any physical tilt develops. If the ZMP approaches the boundary of the geometric support polygon, the controller executes a proactive anchoring maneuver. By commanding immediate feedforward torque allocation and targeted asymmetric bidirectional (reverse) thrust, this maneuver artificially amplifies the contact normal force, thereby safely pinning the chassis to the incline.

Executing this ZMP-guided anchoring requires an adaptive online controller capable of surviving the severe, unmodeled aerodynamic nonlinearities encountered before impact, but crucially without the phase delay inherent to recurrent time-series processing. To explicitly suppress dynamic discrepancies without wasting computational cycles on conventional parameter estimation, which is impractical given the strict Persistency of Excitation (PE) conditions required during millisecond impact transients, Model Reference Adaptive Control (MRAC) is preferred. Specifically, the MRAC Neural Network (MRAC-NN) formulation via Feedback Linearization [10] is theoretically ideal for handling highly coupled aerodynamic disturbances. Unlike standard backpropagation networks requiring offline pre-training and computationally heavy Jacobian calculations, Lewis's MRAC-NN utilizes derivative-free, $\mathcal{O}(1)$ Hebbian weight update rules. This allows continuous learning within a 1000 Hz inner loop while providing mathematically provable Lyapunov stability, ensuring that the network rejects unstructured disturbances without risking algorithmic divergence. Driven by these distinct computational advantages, this study represents the first application of the derivative-free Lewis MRAC-NN framework to the rigid-impact stabilization of an underactuated quadrotor.

Despite its computational simplicity and theoretical power, the classical Lewis framework cannot be directly applied to this problem. The original methodology was explicitly formulated for fully actuated systems (e.g., robotic manipulators and thermal processes) operating strictly within continuous error manifolds. Furthermore, even within a cascaded architecture, standard neuro-adaptive laws degrade rapidly under non-differentiable touchdown constraints. When a quadrotor strikes an incline, it experiences sudden kinematic arrest and sharp derivative spikes. Classical MRAC processes this unilateral physical constraint as a persistent tracking error, integrating its adaptive weights indefinitely to reduce the unresolvable error. Standard static leakage modifications fail to contain these escalating gradients, leading to severe parameter wind-up, high-frequency actuator chattering, and eventual vehicle destabilization.

Main Contributions

To resolve the theoretical gaps between predictive tracking, physical shock absorbers, and classical neuro-adaptive laws, this paper frames inclined touchdown as a real-time, ZMP-bounded robust control task. The main contributions are:

- **Impact-Resilient Underactuated MRAC-NN:** We restructure the computationally efficient, fully actuated Lewis framework [10] into a cascade architecture for controlling a 6-DoF underactuated aerial vehicle. To align with rigorous robust-adaptive control taxonomy, we explicitly decouple the system's unmodeled dynamics into continuous model uncertainties and discontinuous external disturbances. The MRAC Neural Network functions exclusively as a universal approximator to actively learn and **compensate for model uncertainties—**specifically *parametric (structured) uncertainties* (e.g., dynamic inertia shifts) and *non-parametric (unstructured) uncertainties* (e.g., aerodynamic ground-effects and asymmetric thrust allocation). Conversely, unlearnable exogenous shocks, primarily the highly transient *external disturbances* (e.g., rigid impact wrenches during touchdown), are explicitly isolated. To prevent parameter wind-up from these impacts, they are robustly **attenuated** via a dynamic bounding-gain expansion, a dynamic leakage modification, and a kinematic reference-state locking protocol. This decoupled architecture mathematically guarantees Uniform Ultimate Boundedness (UUB) by ensuring the network adapts only to learnable uncertainties while dynamic robust gains suppress unlearnable shocks.
- **Active Anchoring Guided by ZMP:** Rather than applying constant blind reverse thrust upon touchdown, we introduce a deterministic ground-anchoring routine governed by real-time ZMP boundaries. The ZMP isolates the exact threshold of rotational instability before physical tilt occurs, translating the imminent rollover threat into feedforward torque allocation. This dynamically allocates active, asymmetric downward forces to augment the contact normal force, ensuring reliable adherence on steep inclines that exceed the limits of natural static friction.
- **LCP-Based Contact Validation:** Moving beyond idealized penalty-based continuous contact models, the proposed architecture is rigorously validated using a Linear Complementarity Problem (LCP) Dantzig solver coupled with the LuGre tribology

model. This dynamic setup validates the active controller's capability to suppress post-impact bouncing and maintain stability despite ESC hardware latency and severe unmodeled contact transients.

2. Materials and Methods

2.1. Quadrotor Dynamics: A Unified Hybrid Formulation

A hybrid Newton–Euler formulation models the quadrotor dynamics during both free-flight and touchdown phases (Figure 1). Translational states (position $\mathbf{p}^e$ and velocity $\mathbf{v}^e$) are expressed in the inertial Earth frame ($\mathcal{F}_e$), whereas rotational states (Euler attitude $\boldsymbol{\eta}$ and angular velocity $\boldsymbol{\omega}^b$) are defined in the Body frame ($\mathcal{F}_b$). The 6-DoF hybrid twist vector is defined as $\boldsymbol{\nu} = [(\mathbf{v}^e)^T, (\boldsymbol{\omega}^b)^T]^T \in \mathbb{R}^6$. Adopting the mapping $\boldsymbol{\omega}^b = \mathbf{T}_{eul}\dot{\boldsymbol{\eta}}$, the equations of motion are:

$$\mathbf{M}\dot{\boldsymbol{\nu}} + \mathbf{C}(\boldsymbol{\nu})\boldsymbol{\nu} + \mathbf{G}(\boldsymbol{\eta}, \boldsymbol{\nu}, \boldsymbol{\Omega}_{rotor}) = \mathbf{T}_h^T \boldsymbol{\beta}_{true} \mathbf{F}_{rotor}^b + \mathbf{J}_c^T \boldsymbol{\lambda} \tag{1}$$

where $\mathbf{M}, \mathbf{C} \in \mathbb{R}^{6\times 6}$ denotes the hybrid mass-inertia and Coriolis matrices. The generalized wrench $\mathbf{G} \in \mathbb{R}^6$ accounts for gravity, aerodynamic effects (BET + parasitic drag), and gyroscopic reactions (detailed in Appendix B). Ground reaction forces during contact are handled separately via the constraint term $\mathbf{J}_c^T \boldsymbol{\lambda}$. During free-flight ($\boldsymbol{\lambda} \equiv \mathbf{0}$), Equation (1) reduces to the standard underactuated quadrotor dynamics.

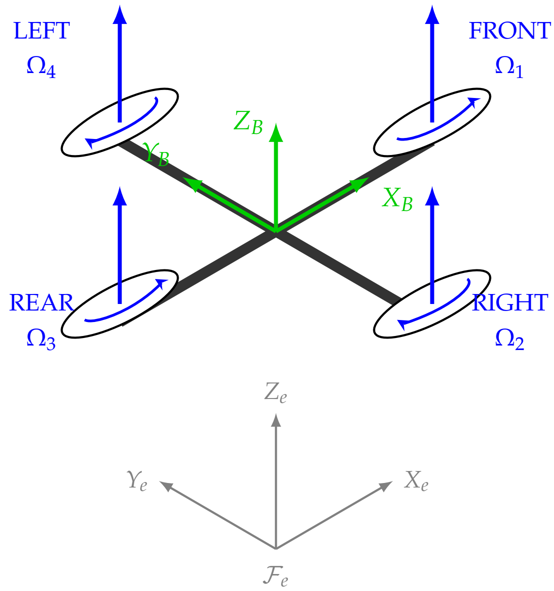

Figure 1. Reference frames and the 6-DoF rigid body dynamics of the quadrotor, illustrating the Earth ($\mathcal{F}_e$) and Body ($\mathcal{F}_b$) frames.

2.1.1. Rotor Actuation and Reverse Thrust

$$F_i = [k_b \mathcal{H}(\Omega_i) - k_{b,rev} \mathcal{H}(-\Omega_i)]\Omega_i^2 \tag{2}$$

where $\mathcal{H}(\cdot)$ denotes the Heaviside step function, and k_b and $k_{b,rev}$ are the forward and reverse thrust coefficients. While Equation (2) provides the nominal control baseline, the simulation computes aerodynamic interactions via BET, introducing state-dependent perturbations to evaluate controller robustness.

2.1.2. Rigid Contact Mechanics and LCP Integration

Upon touchdown ($\boldsymbol{\lambda} \neq \mathbf{0}$), rigid impacts are modeled using acceleration-level contact constraints [11]. Let $\Phi_n(\mathbf{x})$ denote the normal distance to the surface. The second time derivative yields the contact acceleration $\mathbf{a}_c = \mathbf{J}_c\dot{\boldsymbol{\nu}} + \dot{\mathbf{J}}_c\boldsymbol{\nu}$. Substituting $\dot{\boldsymbol{\nu}}$ from (1) reformulates the constraint as a Linear Complementarity Problem (LCP):

$$\begin{aligned}\mathbf{a}_c &= \mathbf{D}_c\boldsymbol{\lambda} \\ &\quad + \mathbf{J}_c\mathbf{M}^{-1}\left(\mathbf{T}_h^T\boldsymbol{\beta}_{true}\mathbf{F}_{rotor}^b - \mathbf{C}\boldsymbol{\nu} - \mathbf{G}\right) + \dot{\mathbf{J}}_c\boldsymbol{\nu},\end{aligned} \tag{3}$$

where $\mathbf{D}_c = \mathbf{J}_c\mathbf{M}^{-1}\mathbf{J}_c^T$ is the Delassus matrix.

Impact Restitution and LuGre Friction: The contact wrench $\boldsymbol{\lambda} = [\lambda_n, \boldsymbol{\lambda}_t^T]^T$ characterizes surface interaction. To maintain continuous-time integration during rigid impacts, normal forces (Signorini–Fichera conditions) and Stribeck stick–slip transitions (LuGre friction [12]) are evaluated simultaneously. An instantaneous kinetic energy dissipation rule applies at collision:

$$\begin{aligned} 0 \le \lambda_{n_i} \perp a_{c,n_i} \ge 0, \quad & v_{cp,n_i}^+ = -e_{res}v_{cp,n_i}^- \\ \dot{\boldsymbol{\zeta}}_i = \mathbf{v}_{cp,t_i} - & \frac{\sigma_{f0}\|\mathbf{v}_{cp,t_i}\|}{g(\mathbf{v}_{cp,t_i})}\boldsymbol{\zeta}_i \\ \boldsymbol{\lambda}_{t_i} = \lambda_{n_i}(\sigma_{f0}\boldsymbol{\zeta}_i & + \sigma_{f1}\dot{\boldsymbol{\zeta}}_i + \sigma_{f2}\mathbf{v}_{cp,t_i}) \end{aligned} \tag{4}$$

Here, the local contact point velocity $\mathbf{v}_{cp,i} = \mathbf{v}^e + \mathbf{R}_b^e(\boldsymbol{\omega}^b \times \mathbf{p}_{leg_i}^b)$ is decomposed into normal ($v_{cp_{n_i}}$) and tangential ($\mathbf{v}_{cp_{t_i}}$) components. The function $g(\mathbf{v}_{cp_{t_i}})$ characterizes the Stribeck effect, and $\boldsymbol{\zeta}_i$ represents internal bristle deflection.

A projected Dantzig algorithm solves the LCP. To evaluate the controller under worst-case rollover conditions, rolling and spinning friction moments are omitted. Surface mechanical stiffness is incorporated into the solver's relaxation parameters, permitting high-frequency contact transients without numerical penetration. To prevent control chattering from sensor noise, each landing gear employs a hysteresis limit switch that triggers the ZMP anchoring logic (Appendix D).

2.2. *Zero Moment Point (ZMP) Formulation for Ground Anchoring*

Widely applied in bipedal robotics [8] and vehicle rollover analysis [13], the Zero Moment Point (ZMP) serves as the primary stability criterion. To prevent a rollover, the ZMP, the planar coordinate at which horizontal overturning moments vanish, must remain within the geometric support polygon ($\mathcal{S}_{diamond}$). Evaluating the ZMP directly in the Body frame ($\mathcal{F}_b$) bypasses trigonometric singularities and eliminates the need for explicit online estimation of the incline angle α. Since IMU specific force and commanded thrust encode the contact wrench direction in the body frame, the formulation yields a computationally lightweight yet geometrically meaningful stability margin for real-time anchoring decisions.

2.2.1. Kinematic ZMP Estimation

Direct measurement of external contact forces during impacts is challenging due to sensor bandwidth limits and ground-effect aerodynamics. Consequently, the ZMP is estimated using onboard IMU data and the transient parameter estimator (Section 2.3.5).

With multi-point contact (three or more legs), the rigid geometry of the landing base forces the body's z-axis to align with the local surface normal on a uniform incline. During the initial single- or dual-leg contact phase, before full alignment is established, IMU-specific force measurements implicitly reflect the direction of the emerging contact wrench; thus, the body-frame ZMP formulation remains a valid directional indicator of tipping tendency without requiring explicit estimation of the incline angle. The MRAC-NN dynamic robustifier and leakage mechanism bound any residual geometric mismatch during this brief transient. Defining the geometric center of the landing base as the origin, the estimated Center of Mass (CoM) position is $\mathbf{p}^b_{com} = [0, 0, h_{com}]^T$, where h_{com} represents the vertical moment arm. While the modeled airframe assumes $h_{com} \approx l$ to simplify the layout, these variables remain mathematically distinct for generality.

The equivalent Ground Reaction Force (GRF) is obtained by subtracting commanded rotor thrusts from mass-scaled IMU specific force measurements (proper acceleration). This yields an estimate of contact-induced ground reaction; aerodynamic effects and model mismatches are compensated for by the MRAC-NN robustifier: $\hat{\mathbf{F}}^b_{grf} = \hat{m}\mathbf{a}^b_{imu} - \mathbf{F}^b_{rotor} = [\hat{F}^b_x, \hat{F}^b_y, \hat{F}^b_z]^T$. Similarly, ground reaction torques are calculated via Euler's equations using gyroscope rates ($\boldsymbol{\omega}^b_{imu}$), derived angular accelerations ($\boldsymbol{\alpha}^b_{imu}$), and commanded rotor torques ($\boldsymbol{\tau}^b_{thrust}$):

$$\begin{aligned} \hat{\boldsymbol{\tau}}^b_{grf} &= \hat{\mathbf{I}}^b_{base}\boldsymbol{\alpha}^b_{imu} + \boldsymbol{\omega}^b_{imu} \times (\hat{\mathbf{I}}^b_{base}\boldsymbol{\omega}^b_{imu}) - \boldsymbol{\tau}^b_{thrust} \\ &= [\hat{\tau}^b_x, \hat{\tau}^b_y, \hat{\tau}^b_z]^T. \end{aligned} \tag{5}$$

The total moment about an arbitrary point on the contact plane ($\hat{\mathbf{p}}^b_{zmp} = [\hat{x}^b_{zmp}, \hat{y}^b_{zmp}, 0]^T$) is expressed as $\hat{\mathbf{M}}^b_{zmp} = (\mathbf{p}^b_{com} - \hat{\mathbf{p}}^b_{zmp}) \times \hat{\mathbf{F}}^b_{grf} + \hat{\boldsymbol{\tau}}^b_{grf}$. By setting the horizontal components of this moment to zero, the estimated ZMP coordinates are analytically extracted. To rigorously validate this estimator against unmodeled dynamics, a precise ground-truth baseline—incorporating the full 3×3 eccentric inertia tensor, kinematic CoM transport, and matched nonlinear filtering—is detailed in Appendix D.

2.2.2. Base Displacement and Polygon Constraints

Since horizontal overturning moments vanish at the ZMP ($\hat{M}^b_x = 0, \hat{M}^b_y = 0$), expanding the cross-product yields the raw ZMP coordinates relative to the geometric center:

$$\hat{x}^b_{raw} = \frac{-h_{com}\hat{F}^b_x - \hat{\tau}^b_y}{\hat{F}^b_z}, \quad \hat{y}^b_{raw} = \frac{-h_{com}\hat{F}^b_y + \hat{\tau}^b_x}{\hat{F}^b_z} \tag{6}$$

To prevent division-by-zero singularities during free-fall or bouncing transients ($\hat{F}^b_z \to 0$), the normal reaction denominator in Equation (6) is constrained to a strictly non-zero minimal threshold. Furthermore, because the numerical differentiation of gyroscope signals introduces high-frequency noise, the raw coordinates ($\hat{x}^b_{raw}, \hat{y}^b_{raw}$) are processed through a discrete-time Fast Nonlinear Tracking Differentiator (FHAN) [14]. This yields dynamically smoothed ZMP positions ($\hat{x}^b_{zmp}, \hat{y}^b_{zmp}$) and their time derivatives ($\dot{\hat{\mathbf{p}}}^b_{zmp}$) for the controller without introducing severe phase lag.

For a "+ (Plus)" configuration, the four landing gears located at $\mathbf{p}^b_i \in \{[\pm l, 0, 0]^T, [0, \pm l, 0]^T\}$ define a diamond-shaped support polygon ($\mathcal{S}_{diamond}$). A rollover occurs if the L_1-norm of the filtered ZMP position ($|\hat{x}^b_{zmp}| + |\hat{y}^b_{zmp}|$) exceeds the arm length (l), indicating that the reaction force vector has exited the support base.

During the anchoring phase, the controller must satisfy the inequality constraint $\hat{\mathbf{p}}^b_{zmp} \in \mathcal{S}_{diamond}$. If contact wrenches ($\mathbf{J}^T_c\boldsymbol{\lambda}$) push the estimated ZMP toward the boundaries of $\mathcal{S}_{diamond}$, the controller commands reverse thrust. As shown in Equation (6), increasing the downward force maximizes the normal reaction magnitude ($\hat{F}^b_z$) in the denominator, which proportionally reduces the ZMP displacement and secures the vehicle to the incline.

2.3. Adaptive Cascade Control Architecture

The control architecture integrates a Model Reference Adaptive Control Neural Network (MRAC-NN) with a Zero Moment Point (ZMP) formulation to maintain stability across unconstrained free-flight and rigid touchdown phases. Based on the hybrid dynamics established in Section 2, the system employs a timescale-separated cascade loop.

An outer loop regulates translational kinematics (e.g., altitude and planar position) in the inertial Earth frame ($\mathcal{F}_e$), while a high-frequency inner loop (MRAC-NN) tracks the resulting attitude references by generating corrective torques in the Body frame ($\mathcal{F}_b$) to compensate for unmodeled uncertainties (Figure 2).

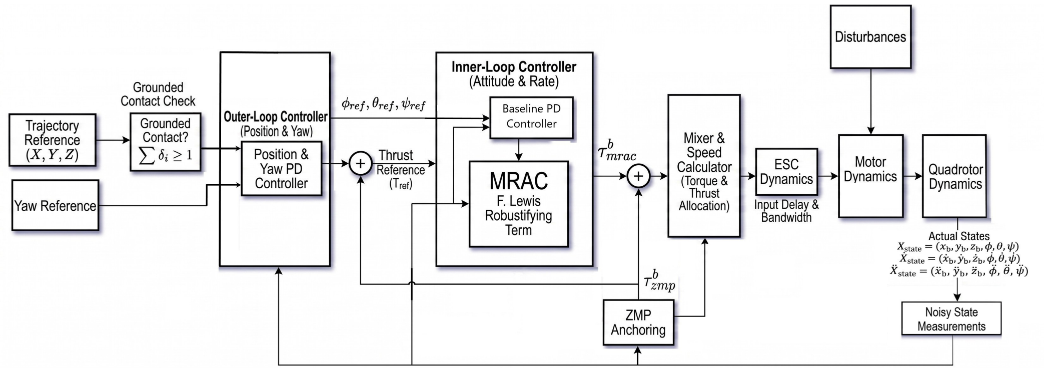

Figure 2. Detailed cascade architecture, illustrating the Outer-Loop (Position/Yaw), Inner-Loop (Attitude/Rate with NNFL), and forward/feedback dynamics.

2.3.1. Outer-Loop: Trajectory and Safety Constraints

The outer loop manages translational trajectory tracking in $\mathcal{F}_e$, decoupling spatial navigation from attitude control. To maintain a consistent 6×1 state dimension across flight phases, the position error dynamics are formulated as:

$$\begin{bmatrix} \mathbf{e}^e_p \\ \mathbf{e}^e_v \end{bmatrix} = \begin{cases} \left[(\mathbf{p}^e_{ref} - \mathbf{p}^e)^T,\ (\mathbf{v}^e_{ref} - \mathbf{v}^e)^T\right]^T, & \text{Airborne} \\ \left[0,\ 0,\ e^e_{p_z},\ 0,\ 0,\ e^e_{v_z}\right]^T, & \text{Contact} \end{cases} \tag{7}$$

During free flight, the controller tracks 3D waypoints. Upon touchdown, ZMP constraints overwrite horizontal position errors, and altitude errors dictate surface adherence. Multiplying the proportional-derivative (PD) efforts by the inner-loop's dynamic neural gain ($K_{nn_{eff}}$) yields the raw Cartesian acceleration demand ($\mathbf{a}^e_{raw}$). Given positive-definite diagonal gain matrices $\mathbf{K}_{p_{out}}$ and $\mathbf{K}_{d_{out}}$, this dynamic coupling functions as a variable stiffness multiplier, automatically adjusting the tracking effort during impacts:

$$\mathbf{a}^e_{raw} = K_{nn_{eff}}\left(\mathbf{K}_{p_{out}}\mathbf{e}^e_p + \mathbf{K}_{d_{out}}\mathbf{e}^e_v\right) \tag{8}$$

A 2nd-order pre-warped Tustin low-pass filter smooths $\mathbf{a}^e_{raw}$ to mitigate sensor noise, yielding $\mathbf{a}^e_{filt}$. The signal progression $\mathbf{a}^e_{raw} \xrightarrow{LPF} \mathbf{a}^e_{filt} \xrightarrow{+g_{est}} \mathbf{a}^e_{des}$ augments the filtered accelerations with a gravity feedforward term and strictly bounds the desired output vector according to the asymmetric physical thrust limits defined in Table 1.

Attitude Extraction and Reference Differentiation: Desired Euler angles (ϕ_{ref}, θ_{ref}) and the baseline vertical thrust command ($u_{1_{out}}$) are extracted [15] by inverting the translational dynamics:

$$\frac{u_{1_{out}}}{\hat{m}}\mathbf{R}^e_b\mathbf{e}_3 = \mathbf{a}^e_{des} + [0, 0, g_{est}]^T \tag{9}$$

where $\mathbf{e}_3 = [0, 0, 1]^T$. Given a yaw reference (ψ_{ref}), the target roll and pitch angles are computed from the third column projection of $\mathbf{R}^e_b$. The thrust magnitude $u_{1_{out}} = \text{sgn}(a^e_{des,z} + g_{est})\hat{m}\|\mathbf{a}^e_{des} + [0, 0, g_{est}]^T\|$ dictates reverse-thrust activation during contact.

The inner-loop MRAC requires higher-order reference derivatives ($\dot{\boldsymbol{\eta}}_{ref}, \ddot{\boldsymbol{\eta}}_{ref}$). Because numerical differentiation amplifies high-frequency noise, a discrete-time Han's Tracking Differentiator (TD) [14] is used to process the raw attitude targets. Upon touchdown, the

TD states are locked ($x_1 = \boldsymbol{\eta}_{est}, x_2 = \mathbf{0}$) to isolate the MRAC from sharp derivative spikes caused by physical impacts.

ZMP Polygon Distance and Smoothing: During the impact phase, the ZMP penetration ($\mathbf{d}^b_{poly}$) is evaluated relative to the support polygon $\mathcal{S}_{diamond}$ (Appendix D). To avoid iterative projection solvers [16], the ZMP coordinates map to the first Cartesian quadrant ($\bar{x} = |\hat{x}^b_{zmp}|, \bar{y} = |\hat{y}^b_{zmp}|$). Based on the nominal arm length l, a $\mathcal{C}^1$ Huber-like penalty function [17] smooths the scalar penetration $s_{raw} = \bar{x} + \bar{y} - l$. Evaluating the cross-couplings ($\delta = 0.5 s_{eff}$) and applying projection tests identify the active Voronoi region. Decoupled magnitudes are then restored to their proper coordinate quadrants using cached signs, yielding the continuous penetration vector $\mathbf{d}^b_{poly}$, which is subsequently differentiated to obtain its rate $\dot{\mathbf{d}}^b_{poly}$.

2.3.2. ZMP-Guided Ground Anchoring and Kinetic State Arbitration

To verify surface interaction without relying on expensive and noise-prone multi-axis force-torque sensors, the landing gear is equipped with low-cost commercial micro limit switches. These physical contact sensors feature a built-in mechanical hysteresis (e.g., 0.75 N actuation threshold and 0.25 N release), which naturally filters out high-frequency aerodynamic or vibrational noise. A discrete positive edge detected from any switch immediately registers a transient *impact* event. This subsequently engages and latches the continuous *contact* mode ($\sum \delta_i \geq 1$) within the flight controller.

Upon verifying this initial ground contact, spatial trajectory tracking is abruptly suspended. Standard cascaded architectures often struggle during rigid impacts because translating an edge-tipping displacement into a corrective attitude reference introduces phase lag, leading to rollover. The logical workflow of the proposed anchoring architecture, designed to overcome these limitations, is illustrated in Figure 3.

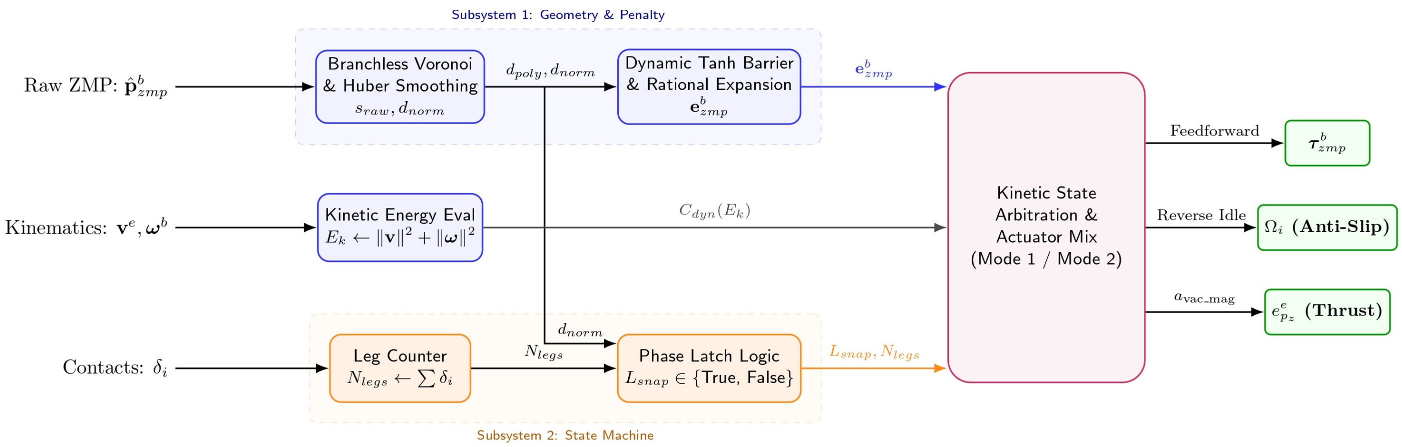

Figure 3. Block diagram of the ZMP anchoring architecture, illustrating the parallel signal flow through the geometric penalty evaluation (Subsystem 1) and the phase latch state machine (Subsystem 2), culminating in the kinetic state arbitration and actuator mix.

To eliminate this phase lag, a strict translational override (referred to as kinematic truncation) forces horizontal position, velocity, and acceleration references to exactly zero ($\mathbf{e}^e_{p_{xy}} = \mathbf{0}, \mathbf{e}^e_{v_{xy}} = \mathbf{0}, \mathbf{a}^e_{ref} = \mathbf{0}$). To prevent reference drift during contact, spatial reference states are clamped to the measured states ($\mathbf{p}^e_{ref} = \mathbf{p}^e$). This action disables lateral navigation entirely, allowing the controller to focus exclusively on rotational alignment and ground adhesion.

Operating independently of the position loop, the ZMP penetration vector ($\mathbf{d}^b_{poly}$) and its derivative ($\dot{\mathbf{d}}^b_{poly}$) are evaluated directly in the body frame. To prevent computational bottlenecks and chattering near the support polygon boundaries, the algorithm employs a branchless Voronoi projection coupled with a $\mathcal{C}^1$-continuous Huber-style penalty. By defining a narrow geometric transition band ($\pm\epsilon$), the controller blends the ZMP error from strict zero into a linear spatial penalty (s_{eff}) without introducing derivative shock waves. This scalar penetration subsequently drives a dynamic barrier multiplier via a bounded hyperbolic tangent function ($T_v = \tanh(\beta_{zmp}[s_{raw} - s_{center}])$). To prevent floating-point unit (FPU) overload during extreme boundary violations, this function incorporates an early-exit hardware saturation guard.

As detailed in Algorithm 1, the proportional-derivative control effort is scaled by a deterministic algebraic rational filter, which continuously expands the torque response as the ZMP approaches the absolute stability limit. Counteracting this tipping displacement requires an immediate orthogonal torque response. Scaling the coupled spatial deviation by the estimated mass ($\hat{m}$) and the arm length (l) computes the corrective feedforward torques ($\boldsymbol{\tau}^b_{zmp}$). To ensure undamped responsiveness, complex gyroscopic smoothing filters are explicitly bypassed. Instead, feedforward torque allocation and reverse-thrust efforts are governed by a robust state-latch and kinetic energy arbitration mechanism.

During the initial unlocked state (*Mode 1: Search and Align*), if only one or two legs are grounded, undamped ZMP feedback actively hinges the vehicle against the slope. Concurrently, instead of applying a constant base vacuum, the algorithm evaluates the instantaneous kinetic energy ($E_k \propto \|\mathbf{v}^e\|^2 + \|\boldsymbol{\omega}^b\|^2$) to scale the downward thrust via a dynamic Kinetic Shock Compressor ($\sqrt[4]{E_k}$). Soft touchdowns inherently bypass this scaling to conserve computational cycles and prevent excessive base-vacuum, whereas high-velocity impacts proactively boost the clutch multiplier (C_{dyn}) to absorb severe rebounds.

Once geometric alignment is achieved (indicated by a negligible ZMP error, $d_{norm} <$ 0.005 m) or solid multi-leg contact is secured ($\sum \delta_i \geq 3$), the Phase Latch State Machine (Subsystem 2) locks into *Mode 2: Impact and Perch*. To prevent mechanical chattering and actuator conflict as the remaining legs make rigid contact with the incline, the ZMP feed-forward torques are instantly deactivated (Kill-Switch engaged, $\boldsymbol{\tau}^b_{zmp} = \mathbf{0}$). Concurrently, vertical velocity tracking is suspended ($e^e_{v_z} = 0$), and an aggressive maximum shock vacuum clutch, further augmented by a discrete snap multiplier (C_{snap}), is injected directly as a proportional vertical position error ($e^e_{p_z}$), forcing absolute ground adhesion to take unconditional precedence. During this transient snap maneuver, while full planar contact is pending ($N_{legs} < 3$), the rotational speeds of the motors corresponding to the already grounded legs ($\delta_i = 1$) are explicitly bypassed from the output mixer by Subsystem 3 and clamped to a constant reverse-thrust command (e.g., $\Omega_{cmd} = 100\,\text{rad/s}$). This aggressive negative thrust forces the grounded contacts to act as a stabilized Reverse Thrust Hinge, pinning the touching legs to the wall and preventing the pivot points from slipping while the ungrounded legs swing into the surface.

Several mechanisms maintain operational stability during the anchoring phase. Primarily, the hardware-level hysteresis of the limit switches acts as a physical debounce filter, ensuring that momentary bounces during a harsh touchdown do not inadvertently disengage the state latch mechanism. Secondarily, the dynamic barrier gain (G_{zmp}) safely bounds the proportional-derivative effort, amplifying the torque response via a hyperbolic tangent function only when the ZMP approaches the geometric boundaries of the support polygon ($\mathcal{S}_{diamond}$). As detailed in Appendix B.3, this formulation assumes minimal planar offsets ($x_{piv}, y_{piv} = 0$); the inner-loop MRAC adaptive robustifier compensates for any minor geometric discrepancies or unmodeled impact transients during execution.

Algorithm 1 Voronoi-Smoothed ZMP Anchoring and Kinetic State Arbitration

Require: Raw ZMP $\hat{\mathbf{p}}^b_{zmp}$, Rate $\dot{\hat{\mathbf{p}}}^b_{zmp}$, Contact array δ_i, Counter c_{solid}, Kinematics $\mathbf{v}^e, \boldsymbol{\omega}^b$
Ensure: Outer-loop override $\mathbf{e}^e_p$, Feedforward torque $\boldsymbol{\tau}^b_{zmp}$, RPM override Ω_i

$N_{legs} \leftarrow \sum \delta_i$
% *1. Subsystem 1: Kinematics Override & Impact Gain Expansion*
$\mathbf{e}^e_{p_{xy}} \leftarrow \mathbf{0}, \quad \mathbf{e}^e_{v_{xy}} \leftarrow \mathbf{0}$ ▷ Force hardware reference to match estimate on ground
$\mathbf{K}_{p_{out}} \leftarrow \alpha_{pd}\mathbf{K}_{p_{out}}, \quad \mathbf{K}_{d_{out}} \leftarrow \alpha_{pd}\mathbf{K}_{d_{out}}$ ▷ Discretely scale up stabilization gains
% *2. Branchless Voronoi Projection & $\mathcal{C}^1$-Continuous Huber Penalty*
$\bar{x} \leftarrow |\hat{x}^b_{zmp}|, \quad \bar{y} \leftarrow |\hat{y}^b_{zmp}|, \quad s_{raw} \leftarrow \bar{x} + \bar{y} - l$ ▷ Evaluate L_1 boundary penetration
if $s_{raw} \leq -\epsilon$ **then**
  $(\mathbf{d}^b_{poly}, \dot{\mathbf{d}}^b_{poly}) \leftarrow (\mathbf{0}, \mathbf{0}), \quad d_{norm} \leftarrow 0$ ▷ Hardware early exit: Stable ZMP inside support polygon
else
  $s_{eff} \leftarrow$ **if** $s_{raw} \geq \epsilon$ **then** s_{raw} **else** $(s_{raw} + \epsilon)^2/(4\epsilon)$ ▷ Apply $\mathcal{C}^1$ Huber smoothing
  $\delta \leftarrow 0.5 s_{eff}$
  $(d_x, d_y) \leftarrow$ **if** $(\bar{x} < \delta)$ **then** $(\bar{x}, s_{eff} - \bar{x})$ **else if** $(\bar{y} < \delta)$ **then** $(s_{eff} - \bar{y}, \bar{y})$ **else** (δ, δ)
  $\mathbf{d}^b_{poly} \leftarrow \left[\operatorname{sgn}(\hat{x}^b_{zmp}) d_x,\ \operatorname{sgn}(\hat{y}^b_{zmp}) d_y,\ 0\right]^T$ ▷ Restore Cartesian quadrants
  Compute $\dot{\mathbf{d}}^b_{poly}$ analogously using analytic directional derivatives
  $d_{norm} \leftarrow \|\mathbf{d}^b_{poly}\|$
  % *3. Dynamic Tanh Barrier & Algebraic Rational Gain Expansion*
  $T_v \leftarrow \tanh\left(\beta_{zmp}(s_{raw} + \epsilon/2)\right)$
  $G_{nom} \leftarrow G_{base} + G_{boost}(1 + T_v)$ ▷ Apply spatial barrier boost
  $\eta_{sat} \leftarrow (G_{nom} \times d_{norm}) \times K_{inv\_comp}$ ▷ Evaluate normalized saturation ratio
  $G_{zmp} \leftarrow G_{nom}\sqrt{1 + \left(\frac{\eta^3_{sat}}{\eta^2_{sat} + 3\alpha_{zmp}}\right)^2}$ ▷ Continuous rational gain expansion
  $\mathbf{e}^b_{zmp} \leftarrow -\left(G_{zmp}\mathbf{d}^b_{poly}\right) - \left(G_{zmp}\dot{\mathbf{d}}^b_{poly} + \dot{G}_{zmp}\mathbf{d}^b_{poly}\right)K_{d_{div}}$
end if
% *4. Subsystem 2: Phase Latch State Machine*
if $c_{solid} \geq 5$ **or** $N_{legs} == 0$ **then**
  $L_{snap} \leftarrow$ **False** ▷ Reset latch on full contact or airborne detection
end if
if $N_{legs} \leq 2$ **and** $d_{norm} < 0.005\,\text{m}$ **then**
  $L_{snap} \leftarrow$ **True** ▷ Geometric alignment strictly achieved, engage latch
end if
% *5. Subsystem 3: Torque Arbitration & Actuator Mix*
if not L_{snap} **and** $N_{legs} \leq 2$ **then**
  $\boldsymbol{\tau}^b_{zmp} \leftarrow \left[e^b_{zmp,y} K_{p_\phi} \hat{m} l,\ e^b_{zmp,x} K_{p_\theta} \hat{m} l,\ 0\right]^T$ ▷ Mode 1: Search & Align. Cross-coupled torque.
  $E_k \leftarrow (\|\mathbf{v}^e\|^2 + \|\boldsymbol{\omega}^b\|^2) K_{kin}$ ▷ Evaluate instantaneous kinetic energy
  $C_{dyn} \leftarrow \min\left(C_{max},\ C_{base} \times \textbf{if } E_k < 1.0 \textbf{ then } 1.0 \textbf{ else } \sqrt[4]{E_k}\right)$
  $a_{vac\_mag} \leftarrow a^-_{max} \times C_{dyn}$ ▷ Dynamic shock absorption
else
  $\boldsymbol{\tau}^b_{zmp} \leftarrow \mathbf{0}$ ▷ Mode 2: Impact & Perch (Torque Kill-Switch)
  if L_{snap} **and** $N_{legs} \leq 2$ **then**
    $a_{vac\_mag} \leftarrow a^-_{max} \times C_{max} \times C_{snap}$ ▷ Aggressive adherence multiplier
    $\Omega_i \leftarrow \Omega_{anti\_slip} \quad \forall i \in \{j \mid \delta_j = 1\}$ ▷ Force active hinge via reverse idle on touching legs
  else
    $a_{vac\_mag} \leftarrow a^-_{max} \times C_{max}$ ▷ Stable full-surface contact
  end if
end if
% *6. Thrust Override via Vertical Error Projection*
$e^e_{p_z} \leftarrow -\left(a_{vac\_mag} + \hat{g}_{est} K_{grav\_scale}\right)\left(K_{p_z} K_{nn_{eff}}\right)^{-1}$ ▷ Negative Z thrust formulation

2.3.3. Inner-Loop: Neural Network Feedback Linearization

Neural Network Feedback Linearization maps the coupled rigid-body dynamics into a linear error manifold. The adaptive update laws are formulated based on the Lewis framework [10], utilizing a robust Hebbian tuning method instead of standard backpropagation. Although stability proofs rely on Taylor series expansions, real-time execution remains derivative-free, avoiding Jacobian calculations to maintain $\mathcal{O}(1)$ computational efficiency.

Composite Error and Actuated Manifold Synthesis: The MRAC evaluates a 6-DoF composite tracking error vector defined as $\mathbf{r}_c = \dot{\mathbf{e}} + \mathbf{\Lambda e} \in \mathbb{R}^{6\times1}$. Preserving the unscaled derivative formulation ($\dot{\mathbf{e}}$) is critical to maintain the skew-symmetry property in the Lyapunov stability analysis. To explicitly embed the physical cascade Proportional-Derivative (PD) gains into this adaptive manifold, the bandwidth matrix is analytically defined as $\mathbf{\Lambda} = \mathbf{K}_d^{-1}\mathbf{K}_p$. Here, $\mathbf{K}_p, \mathbf{K}_d \in \mathbb{R}^{6\times6}$ are strictly positive-definite diagonal matrices structured to encompass the full 6-DoF tracking feedback (e.g., $\mathbf{K}_p = \mathrm{diag}(K_{p_x}, K_{p_y}, K_{p_z}, K_{p_\phi}, K_{p_\theta}, K_{p_\psi})$). Given this multi-dimensional positive-definite constraint, the differential equation $\mathbf{r}_c = \mathbf{0}$ enforces exponentially stable error dynamics across all operational axes. Bounding this composite error guarantees that the spatial trajectory deviations are Uniformly Ultimately Bounded (UUB): $\|\mathbf{e}\| \le b_e$, where $\mathbf{e} = [(\mathbf{p}^e_{ref} - \mathbf{p}^e)^T,\ (\boldsymbol{\eta}_{ref} - \boldsymbol{\eta})^T]^T \in \mathbb{R}^{6\times1}$.

To isolate the actively controlled states (altitude and attitude) from the underactuated planar drift (X and Y axes), the actuated error sub-manifold is defined as $\bar{\mathbf{r}}_c = \mathbf{P}\mathbf{r}_c$, utilizing the projection mask $\mathbf{P} = \mathrm{diag}(0, 0, 1, 1, 1, 1)$. To drive neural activation, a 30-dimensional state vector $\mathbf{X}_{nn} \in \mathbb{R}^{30\times1}$ incorporates kinematic errors, operational references, and temporal derivatives:

$$\mathbf{X}_{nn} = \left[\mathbf{e}^T,\ \dot{\mathbf{e}}^T,\ (\mathbf{p}^e_{ref})^T,\ \boldsymbol{\eta}^T_{ref},\ (\mathbf{v}^e_{ref})^T,\ (\dot{\boldsymbol{\eta}}_{ref})^T,\ (\mathbf{a}^e_{ref})^T,\ (\ddot{\boldsymbol{\eta}}_{ref})^T\right]^T \tag{10}$$

Control Synthesis and Gain Expansion: Let $\hat{\mathbf{\Phi}}_1 = \tanh(\hat{\mathbf{V}}^T\mathbf{X}_{nn})$ define the hidden layer activation vector. To suppress unstructured disturbances, a robustifying signal ($\mathbf{u}_{rob}$) and a dynamic adaptive gain ($K_{nn_{eff}}$) are calculated using the actuated error manifold ($\bar{\mathbf{r}}_c$). Defining the operational envelope limit as $M_{safe} = K_{safe} - Z_B$ (where K_{safe} bounds the maximum allowable adaptive gain), the tuning equations evaluate as [10]:

$$\mathbf{u}_{rob} = -K_z\bar{\mathbf{r}}_c\left(\|\hat{\mathbf{Z}}\|_F + Z_B\right) \tag{11}$$

$$K_{nn_{eff}} = Z_B + M_{safe}\tanh\left(\frac{\hat{c}_1 + \eta_s\tanh(k_s\|\bar{\mathbf{r}}_c\|)}{M_{safe}}\right) \tag{12}$$

where $\|\hat{\mathbf{Z}}\|_F = \sqrt{\|\hat{\mathbf{W}}\|_F^2 + \|\hat{\mathbf{V}}\|_F^2}$ and Z_B denote the theoretical boundary parameters.

The neural approximation ($\hat{\mathbf{W}}^T\hat{\mathbf{\Phi}}_1$) forms the base pseudo-control effort. Rather than applying a uniform scalar gain that disregards the distinct inertial properties of the quadrotor's rotational axes, the composite error is first scaled by the diagonal derivative gain matrix ($\mathbf{K}_d$). Because expanding this structure algebraically yields $\mathbf{K}_d\bar{\mathbf{r}}_c = \mathbf{K}_d\dot{\bar{\mathbf{e}}} + \mathbf{K}_p\bar{\mathbf{e}}$, this step mathematically recovers the exact decoupled linear PD efforts. Consequently, the scalar dynamic neural gain ($K_{nn_{eff}}$) acts solely as a transient multiplier over this structured baseline, amplifying adherence forces during impacts without violating passivity:

$$\mathbf{u}_{base} = K_{nn_{eff}}\mathbf{K}_d\bar{\mathbf{r}}_c + \mathbf{P}\hat{\mathbf{W}}^T\hat{\mathbf{\Phi}}_1 - \mathbf{u}_{rob} \tag{13}$$

Torque Assembly and ZMP Superposition: Since $\mathbf{u}_{base}$ defines commanded pseudo-accelerations in the inertial Euler space, its rotational components ($\mathbf{u}_{base_{4:6}}$) must be projected into the body frame ($\mathcal{F}_b$) and scaled by the estimated inertia tensor to recover the adaptive torque. Incorporating the Euler rate Jacobian ($\mathbf{T}_{eul}$) and a feedforward kinematic correction vector $\mathbf{C}_{kin} = \dot{\mathbf{T}}_{eul}\dot{\boldsymbol{\eta}}_{ref}$ to address transient orientation shifts, the MRAC torque is computed directly as:

$$\boldsymbol{\tau}^b_{mrac} = \hat{\mathbf{I}}^b_{base}\left(\mathbf{T}_{eul}\mathbf{u}_{base_{4:6}} + \mathbf{C}_{kin}\right) \tag{14}$$

The final control law combines the total vertical thrust (u_1) and the body-frame control torques ($\boldsymbol{\tau}^b_{total}$). While the outer loop governs the primary altitude trajectory via $u_{1_{out}}$, the inner loop contributes its mass-scaled adaptive Z-axis compensation (u_{nn_3}, representing the third element of $\mathbf{u}_{base}$). The commanded thrust is bounded by physical limits, and the total torque superimposes the feedforward ZMP response ($\boldsymbol{\tau}^b_{zmp}$) directly onto the adaptive MRAC output:

$$u_1 = u_{1_{out}} + \hat{m}u_{nn_3} \tag{15}$$

$$\boldsymbol{\tau}^b_{total} = \boldsymbol{\tau}^b_{mrac} + \boldsymbol{\tau}^b_{zmp} \tag{16}$$

During multi-point contact, the spatial tracking error is driven to zero ($\bar{\mathbf{r}}_c \approx \mathbf{0}$). This naturally attenuates the proportional-derivative effort and freezes the adaptive neural weights at their steady-state values ($\dot{\hat{\mathbf{W}}} \approx \mathbf{0}$), ensuring the transient feedforward ZMP torque unconditionally dominates the adherence phase without adversarial actuator conflict.

Decoupled Hebbian Learning and Weight Adaptation: To manage disparate timescales (e.g., fast roll/pitch control versus slower yaw dynamics), the standard learning rate scalars are expanded into diagonal matrices $\boldsymbol{\Gamma}_W$ and $\boldsymbol{\Gamma}_V$. Derivative-free execution bypasses the backpropagation chain rule. Using a dynamic leakage scalar ($\kappa = \hat{c}_1/Z_B$), the continuous-time weight adaptation equations evaluate the actuated manifold to bound parameter growth:

$$\dot{\hat{\mathbf{W}}} = \hat{\boldsymbol{\Phi}}_1(\boldsymbol{\Gamma}_W\hat{\mathbf{M}}_x\bar{\mathbf{r}}_c)^T - \kappa\|\bar{\mathbf{r}}_c\|\hat{\mathbf{W}}\boldsymbol{\Gamma}_W \tag{17}$$

$$\dot{\hat{\mathbf{V}}} = \boldsymbol{\Gamma}_V\|\bar{\mathbf{r}}_c\|\mathbf{X}_{nn}\hat{\boldsymbol{\Phi}}_1^T - \kappa\|\bar{\mathbf{r}}_c\|\boldsymbol{\Gamma}_V\hat{\mathbf{V}} \tag{18}$$

where $\hat{\mathbf{M}}_x$ is the estimated generalized mass-inertia matrix (detailed in Appendix A). According to the foundational MRAC bounding lemmas [10], the norm of the neural network input vector ($\mathbf{X}_{nn}$) and the subsequent reconstruction errors are bounded by theoretical constants. Because analytically computing the state-independent bound (c_1) is impossible prior to flight due to unknown initial errors and impact transients, the aggregate upper bound is treated strictly as an unknown constant. Consequently, an adaptive boundary scalar $\hat{c}_1$ is integrated online to dynamically estimate this parameter via the following differential law:

$$\dot{\hat{c}}_1 = \gamma\|\bar{\mathbf{r}}_c\| - \sigma_{c1}\hat{c}_1 \tag{19}$$

where $\gamma > 0$ is the adaptation rate and $\sigma_{c1} > 0$ is the robust leakage parameter. To prevent numerical drift during high-frequency impact events, the adaptive weight updates and Equation (19) are integrated using a 2nd-order Exponentially Time-Differenced (ETD2) scheme [18].

2.3.4. Dynamic Control Allocation

The resulting control effort $\mathbf{U} = [u_1, (\boldsymbol{\tau}^b_{total})^T]^T \in \mathbb{R}^4$ is mapped to individual rotor forces $\mathbf{F} \in \mathbb{R}^4$ via an active allocation matrix $\boldsymbol{\beta}_{act}(\mathbf{s}) \in \mathbb{R}^{4\times4}$. This matrix is constructed by extracting the bottom four rows of the nominal geometry ($\boldsymbol{\beta}_{nom}$, defined in Appendix B.3), inherently discarding the unactuated planar dimensions, and dynamically replacing the baseline yaw parameters with the rotational sign vector $\mathbf{s} = \mathrm{sgn}(\boldsymbol{\Omega}_{rotor})$. The unconstrained thrust commands are initially computed via direct analytical inversion, $\mathbf{F} = \boldsymbol{\beta}^{-1}_{act}(\mathbf{s})\mathbf{U}$. If these calculated forces violate bidirectional actuator limits ($f_{\min} \leq F_i \leq f_{\max}$), an active-sign prediction sequence attempts to analytically bound the solution. If this projection remains infeasible, the system falls back to an iterative Quadratic Programming (QP) solver (e.g., OSQP [19]) to compute the optimal bounded thrusts. Crucially, during the active snap phase ($L_{snap} =$ **True** and $N_{legs} \leq 2$), this analytical allocation for the grounded rotors is entirely bypassed. Their velocity commands are strictly locked to the reverse-idle threshold to enforce the mechanical hinge, ensuring uninterrupted physical anchoring isolated from transient IMU spikes.

2.3.5. Initial Parameter Estimation for ZMP Anchoring

Accurate identification of the vehicle mass ($\hat{m}$) and inertia tensor ($\hat{\mathbf{I}}^b_{base}$) prevents spatial scaling errors during Phase 2 ZMP anchoring. To bypass computationally expensive inverse-dynamics algorithms, parameter estimation is divided into two routines executed during the initial takeoff.

First, mass adaptation is restricted to the vertical axis. To prevent noise-induced drift, a zero-leakage barrier activates: the leakage penalty is disabled while $\hat{m}$ remains within a predefined safe envelope. Boundary violations trigger an exponential penalty to confine the estimate, subsequently processed via a 20 Hz Exponential Moving Average (EMA).

Second, an Extended Kalman Filter (EKF) estimates the roll and pitch rotational inertias [20] ($\hat{I}_{xx}$, $\hat{I}_{yy}$), while the yaw inertia ($\hat{I}_{zz}$) is scaled relative to the active mass ratio ($\hat{m}/m_{base}$). Continuous adaptation during aggressive impacts can corrupt covariance matrices. Therefore, both estimators are constrained to an initial temporal window. A sigmoid function fades learning rates to zero shortly after takeoff. This isolates parameter estimation from the landing phase, locking the parameters to provide stable scaling factors for ZMP torque calculations.

2.4. Stability Analysis and Closed-Loop Error Dynamics

The control objective is to ensure that the composite tracking error converges to a bounded region while keeping the neural network weights bounded to prevent parameter drift.

2.4.1. Closed-Loop Error Formulation

Spatial offsets ($\mathbf{p}^b_{com}, \mathbf{p}^b_{piv} \neq \mathbf{0}$) make the physical plant ($\mathbf{M}, \mathbf{V}$, defined in Appendix A) asymmetric. However, standard MRAC assumes a symmetric structure ($\mathbf{p}^b_{com} = \mathbf{p}^b_{piv} = \mathbf{0}$), defining the nominal mass-inertia matrix ($\mathbf{M}_{nom}$) as block-diagonal. Projected into generalized coordinates $\mathbf{x} = [(\mathbf{p}^e)^T, \boldsymbol{\eta}^T]^T$ using the kinematic mapping $\boldsymbol{\nu} = \mathbf{J}_q\dot{\mathbf{x}}$ (with $\mathbf{J}_q$ defined in Appendix C), the open-loop dynamics are:

$$\mathbf{M}_x\dot{\mathbf{r}}_c = -\mathbf{C}_x\mathbf{r}_c - \mathbf{u}_x + \mathbf{F}_{ideal} + \boldsymbol{\Delta}_{unc}(\mathbf{x}) - \mathbf{J}^T_{c,x}\boldsymbol{\lambda} + \mathbf{d}_{ext}(t) \tag{20}$$

where $\mathbf{M}_x = \mathbf{J}^T_q\mathbf{M}_{nom}\mathbf{J}_q$ and $\mathbf{C}_x = \mathbf{J}^T_q(\mathbf{M}_{nom}\dot{\mathbf{J}}_q + \mathbf{C}_{nom}\mathbf{J}_q)$. Since $\mathbf{M}_{nom}$ is symmetric, the matrix ($\dot{\mathbf{M}}_x - 2\mathbf{C}_x$) remains skew-symmetric. Here, the theoretical formulation assumes full actuation to define an idealized 6-DoF nominal feedforward wrench $\mathbf{F}_{ideal} \in \mathbb{R}^{6\times1}$, and $\mathbf{J}_{c,x}$ is the generalized contact Jacobian mapped from the hybrid manifold (detailed in Appendix C). Correspondingly, the generalized control effort demanded by the baseline adaptive framework is defined as $\mathbf{u}_x \in \mathbb{R}^{6\times1}$.

Unlike standard robust formulations that blindly lump all unknown dynamics into a generic disturbance term, we explicitly decouple the system discrepancies based on their physical and spectral characteristics:

- **Model Uncertainties ($\boldsymbol{\Delta}_{unc}(\mathbf{x})$):** This encapsulates state-dependent, continuous discrepancies. It includes the *parametric (structured) uncertainty* caused by the dynamic inertial mismatch $\Delta_{dyn} = (\mathbf{M}_{nom} - \mathbf{M})\dot{\boldsymbol{\nu}} + (\mathbf{C}_{nom} - \mathbf{C})\boldsymbol{\nu}$, and the *non-parametric (unstructured) uncertainty* arising from structural control allocation mismatch ($\Delta\boldsymbol{\beta}$) mapping the 6-DoF demand to the 4-DoF physical quadrotor, as well as unmodeled continuous aerodynamic wall-effects.

- **External Disturbances ($\mathbf{d}_{ext}(t)$ and $\boldsymbol{\lambda}$):** This encapsulates highly transient, stochastic, and explicitly unlearnable exogenous variables. This includes stochastic environmental noise ($\mathbf{d}_{ext}$) and the discontinuous rigid impact wrenches ($-\mathbf{J}_{c,x}^T\boldsymbol{\lambda}$) encountered upon touchdown.

The strict theoretical purpose of the MRAC Neural Network is to function as a universal approximator to actively learn and **compensate** for the continuous model uncertainties. Thus, using the estimated mass-inertia matrix $\hat{\mathbf{M}}_x$, the target coupled nonlinear function to be parameterized by the NN is physically redefined as $\mathbf{f}(\mathbf{x}) = \hat{\mathbf{M}}_x^{-1}\big(\mathbf{F}_{ideal} + \boldsymbol{\Delta}_{unc}(\mathbf{x})\big) \in \mathbb{R}^6$. Substituting the control effort $\mathbf{u}_x = \hat{\mathbf{M}}_x\mathbf{u}_{base}$ and defining the functional approximation error as $\tilde{\mathbf{f}} = \tilde{\mathbf{W}}^T\hat{\boldsymbol{\Phi}}_1 + \mathbf{w}_{dist} + \boldsymbol{\varepsilon}$, the open-loop dynamics are projected onto the actuated sub-manifold ($\bar{\mathbf{r}}_c = \mathbf{P}\mathbf{r}_c$). Because the spatial projection mask ($\mathbf{P}$) mathematically commutes with the nominal block-diagonal matrices ($\mathbf{P}\mathbf{M}_x = \mathbf{M}_x\mathbf{P}$ and $\mathbf{P}\mathbf{C}_x = \mathbf{C}_x\mathbf{P}$), the closed-loop tracking error dynamics simplify elegantly to:

$$\mathbf{M}_x\dot{\bar{\mathbf{r}}}_c = -\mathbf{C}_x\bar{\mathbf{r}}_c - \hat{\mathbf{M}}_x K_{nn_{eff}}\mathbf{K}_d\bar{\mathbf{r}}_c + \hat{\mathbf{M}}_x\mathbf{P}\tilde{\mathbf{W}}^T\hat{\boldsymbol{\Phi}}_1 + \mathbf{D}_{ext} + \hat{\mathbf{M}}_x\mathbf{u}_{rob} \tag{21}$$

Because the adaptive neural network explicitly cancels the continuous model uncertainties ($\hat{\mathbf{M}}_x\mathbf{P}\tilde{\mathbf{W}}^T\hat{\boldsymbol{\Phi}}_1 \approx \mathbf{P}\boldsymbol{\Delta}_{unc}$), the residual projected vector is redefined as $\mathbf{D}_{ext} = \hat{\mathbf{M}}_x\mathbf{d}_{nn} - \mathbf{P}\mathbf{J}_{c,x}^T\boldsymbol{\lambda} + \mathbf{P}\mathbf{d}_{ext}(t)$. Crucially, this lumped residual disturbance now *exclusively* encapsulates the unlearnable external impact disturbances, environmental noise, and the bounded network reconstruction error.

Extending the standard analytical bounding principles of robust control to encompass finite-energy impact wrenches, this strictly exogenous residual is conservatively assumed to satisfy the linearly parameterized upper bound: $\|\mathbf{D}_{ext}\| \leq c_1 + c_2\|\bar{\mathbf{r}}_c\|$. Because discontinuous external shocks ($\boldsymbol{\lambda}$) cannot be parameterized by continuous basis functions, they must be robustly **attenuated**. The state-dependent component ($c_2\|\bar{\mathbf{r}}_c\|$)—which physically scales with the kinetic collision energy—is subsequently absorbed by the robust derivative gain matrix $\mathbf{K}_d$ during stabilization. This leaves c_1 as the incomputable state-independent upper bound for stochastic external noise. Because the inner-layer weight update uses a derivative-free rule, the uncanceled Taylor series errors are bounded and absorbed into the network reconstruction error. This error is defined strictly within the projected space as $\mathbf{d}_{nn} = \mathbf{P}(\mathbf{w}_{dist} + \boldsymbol{\varepsilon} + \boldsymbol{\epsilon}_V)$, where $\mathbf{w}_{dist}$ is the ideal weight approximation error, $\boldsymbol{\varepsilon}$ is the bounded activation error, and $\boldsymbol{\epsilon}_V$ contains the bounded uncanceled terms from the hidden layer.

2.4.2. UUB Analysis and Dynamic Gain Integration

Assumption 1. ***(Hybrid Flow-Jump Consistency):*** *To rigorously align the continuous-time adaptive laws with the non-differentiable rigid-impact dynamics and discontinuous state-machine transitions, the operational domain is analyzed across continuous flow intervals separated by discrete jump events.*

1. ***Continuous Flow (Inter-Impact Intervals):*** *During free-flight ($\boldsymbol{\lambda} \equiv \mathbf{0}$) and sustained continuous contact (where the normal acceleration $a_{c_n} = 0$ and $\boldsymbol{\lambda}$ is finite), the lumped spatial disturbance satisfies the linearly parameterized bound $\|\mathbf{D}_{ext}\| \leq c_1 + c_2\|\bar{\mathbf{r}}_c\|$. The continuous-time Lyapunov analysis applies strictly over these flow intervals.*
2. ***Discrete Jumps (Rigid Collisions and Switching):*** *At the infinitesimal instant of a rigid touchdown (t_k), the contact wrench constitutes a non-differentiable impulse ($\boldsymbol{\lambda} \to$*

∞). Due to the inelastic contact restitution ($e_{res} < 1$) and the dynamic Kinetic Shock Compressor (C_{dyn}), the collision strictly dissipates generalized kinetic energy. Furthermore, the discontinuous state-machine transitions (e.g., kinematic reference locking $\mathbf{e}^{e}_{p_{xy}} \to \mathbf{0}$) instantaneously project the tracking error to a lower-magnitude state. Therefore, it is assumed that the Lyapunov candidate function experiences a non-positive discrete jump across the impact: $V(t_k^+) \leq V(t_k^-)$.

Theorem 1. *Consider the closed-loop tracking error dynamics given by Equation (21) subject to the lumped spatial disturbance $\mathbf{D}_{total}$. If the control effort is applied alongside the continuous-time adaptation laws for the neural weights and the dynamic bounding scalar, and the baseline stabilization gain satisfies $m_{\min}k_{d,\min} > 0$, then the composite tracking error $\bar{\mathbf{r}}_c$ and the parameter estimation error $\tilde{\mathbf{Z}}$ are Uniformly Ultimately Bounded (UUB).*

Proof. To establish UUB [21], we define a positive-definite Lyapunov candidate using the kinetic tracking energy and parameter estimation errors ($\tilde{\mathbf{Z}} = \text{blkdiag}(\tilde{\mathbf{W}}, \tilde{\mathbf{V}})$, $\tilde{c}_1 = c_1 - \hat{c}_1$):

$$V = \frac{1}{2}\bar{\mathbf{r}}_c^T \mathbf{M}_x \bar{\mathbf{r}}_c + \frac{1}{2}\text{tr}\{\tilde{\mathbf{W}}\mathbf{\Gamma}_W^{-1}\tilde{\mathbf{W}}^T\} + \frac{1}{2}\text{tr}\{\tilde{\mathbf{V}}^T\mathbf{\Gamma}_V^{-1}\tilde{\mathbf{V}}\} + \frac{1}{2\gamma}\tilde{c}_1^2 \tag{22}$$

Since both the nominal mass-inertia matrix $\mathbf{M}_x$ and the Coriolis matrix $\mathbf{C}_x$ commute with the projection mask ($\mathbf{P}^T\mathbf{M}_x\mathbf{P} = \mathbf{P}\mathbf{M}_x$, $\mathbf{P}^T\mathbf{C}_x\mathbf{P} = \mathbf{P}\mathbf{C}_x$) along the decoupled manifold, the fundamental skew-symmetry property $\bar{\mathbf{r}}_c^T(\frac{1}{2}\dot{\mathbf{M}}_x - \mathbf{C}_x)\bar{\mathbf{r}}_c = 0$ is preserved within the projected space. Differentiating Equation (22) and substituting the adaptation laws cancels the parameter estimation error $\bar{\mathbf{r}}_c^T\hat{\mathbf{M}}_x\tilde{\mathbf{W}}^T\hat{\mathbf{\Phi}}_1$.

Substituting the closed-loop tracking error dynamics Equation (21) into the time derivative of the Lyapunov function explicitly introduces the proportional-derivative control effort into the energy dissipation. Thus, the derivative gain matrix acts as the primary robust stabilizer. Since $\hat{\mathbf{M}}_x$ is symmetric positive-definite and $\mathbf{K}_d$ is strictly diagonal positive-definite, the quadratic form $\bar{\mathbf{r}}_c^T\hat{\mathbf{M}}_x\mathbf{K}_d\bar{\mathbf{r}}_c$ strictly dissipates energy. Assuming the baseline stabilization gain is tuned sufficiently large to absorb the state-dependent dynamic disturbance component c_2, the effective minimum eigenvalue of the symmetric part of the stabilization matrix is defined via the Rayleigh quotient theorem as:

$$\lambda_{min}\left(\tfrac{1}{2}(\hat{\mathbf{M}}_x\mathbf{K}_d + \mathbf{K}_d\hat{\mathbf{M}}_x)\right) - c_2 \equiv m_{\min}k_{d,\min} > 0. \tag{23}$$

Consequently, the net stabilization effort bounding the coupled error dynamics satisfies the following inequality:

$$\bar{\mathbf{r}}_c^T\hat{\mathbf{M}}_x\mathbf{K}_d\bar{\mathbf{r}}_c - c_2\|\bar{\mathbf{r}}_c\|^2 \geq m_{\min}k_{d,\min}\|\bar{\mathbf{r}}_c\|^2. \tag{24}$$

Furthermore, since the robustifying signal yields a negative semi-definite dissipation term ($\bar{\mathbf{r}}_c^T\hat{\mathbf{M}}_x\mathbf{u}_{rob} \leq 0$), it strictly absorbs unstructured residual energy and is conservatively omitted from the upper bound. Using the remaining state-independent bound c_1, the Lyapunov derivative evaluates as:

$$\dot{V} \leq -m_{\min}k_{d,\min}K_{nn_{eff}}\|\bar{\mathbf{r}}_c\|^2 + \|\bar{\mathbf{r}}_c\|c_1 + \kappa\|\bar{\mathbf{r}}_c\|\text{tr}\{\tilde{\mathbf{Z}}\hat{\mathbf{Z}}^T\} - \frac{1}{\gamma}\tilde{c}_1\dot{\hat{c}}_1 \tag{25}$$

Substituting the adaptation law $\dot{\hat{c}}_1 = \gamma\|\bar{\mathbf{r}}_c\| - \sigma_{c1}\hat{c}_1$ decomposes the bounding term as $-\frac{1}{\gamma}\tilde{c}_1\dot{\hat{c}}_1 = -c_1\|\bar{\mathbf{r}}_c\| + \hat{c}_1\|\bar{\mathbf{r}}_c\| + \frac{\sigma_{c1}}{\gamma}\tilde{c}_1\hat{c}_1$. The negative component strictly cancels the physical disturbance bound $c_1\|\bar{\mathbf{r}}_c\|$, leaving a bounded residual $\hat{c}_1\|\bar{\mathbf{r}}_c\| \leq \hat{c}_{1,max}\|\bar{\mathbf{r}}_c\|$ that is subsequently absorbed into D_R. To simplify the trace expansions, we apply the dynamic leakage rule $\kappa = \hat{c}_1/Z_B$ and the standard robust adaptive trace inequality [22] $\text{tr}\{\tilde{\mathbf{Z}}\hat{\mathbf{Z}}^T\} \leq$

$\|\tilde{\mathbf{Z}}\|_F Z_B - \|\tilde{\mathbf{Z}}\|_F^2$. Completing the square for both the parameter matrices and the bounding scalar dynamics ($\dot{\hat{c}}_1 = \gamma\|\bar{\mathbf{r}}_c\| - \sigma_{c1}\hat{c}_1$) yields the intermediate dissipation inequality:

$$\begin{aligned}\dot{V} \leq & - m_{\min} k_{d,\min} K_{nn_{eff}} \|\bar{\mathbf{r}}_c\|^2 + D_R \|\bar{\mathbf{r}}_c\| \\ & - \kappa \|\bar{\mathbf{r}}_c\| \left(\|\tilde{\mathbf{Z}}\|_F - C_Z\right)^2 - \frac{\sigma_{c1}}{\gamma}\left(\tilde{c}_1 - C_c\right)^2 + D_c\end{aligned} \tag{26}$$

where $D_R \triangleq \hat{c}_{1,max} + \kappa_{max}\frac{Z_B^2}{4}$ bounds the total residual disturbance, integrating the uncanceled parameter estimation remnant (with $\kappa_{max} = \hat{c}_{1,max}/Z_B$ defined by a preset software saturation limit $\hat{c}_{1,max}$). The unknown theoretical constants are $C_Z = Z_B/2$, $C_c = c_1/2$, and $D_c = \sigma_{c1} c_1^2 / 4\gamma$.

In standard MRAC, a static bounding gain ($K_{nn_{eff}} \equiv K_{min}$) causes the ultimate tracking bound to diverge during impact transients due to the rapid unconstrained growth of the lumped spatial disturbance D_R. To counteract this structural vulnerability, the dynamically expanding neural gain $K_{nn_{eff}}(\bar{\mathbf{r}}_c, \hat{c}_1)$—formulated earlier to track error peaks—is substituted directly into the intermediate Lyapunov inequality. Completing the square for the primary tracking term isolates the bounded residual disturbance D_{UUB} and yields the decoupled Lyapunov derivative:

$$\begin{aligned}\dot{V} \leq & - m_{\min} k_{d,\min} K_{nn_{eff}}(\cdot) \left(\|\bar{\mathbf{r}}_c\| - \frac{D_R}{2 m_{\min} k_{d,\min} K_{nn_{eff}}(\cdot)}\right)^2 \\ & - \kappa \|\bar{\mathbf{r}}_c\| \left(\|\tilde{\mathbf{Z}}\|_F - C_Z\right)^2 - \frac{\sigma_{c1}}{\gamma}\left(\tilde{c}_1 - C_c\right)^2 + \underbrace{D_c + \frac{D_R^2}{4 m_{\min} k_{d,\min} K_{nn_{eff}}(\cdot)}}_{D_{UUB}}\end{aligned} \tag{27}$$

During a rigid impact ($\lambda \to \infty$), the transient disturbance quadratically inflates the D_R^2 numerator. However, the concurrent growth of the tracking error ($\|\bar{\mathbf{r}}_c\|$) and the bounding scalar ($\hat{c}_1$) drives the nested hyperbolic tangent of the dynamic gain toward unity. At the saturation boundary $\mathcal{B}_{sat} \triangleq \{\|\bar{\mathbf{r}}_c\| = r_{peak}, \hat{c}_1 = \hat{c}_{1,max}\}$, the dynamic gain reaches its hardware-defined operational maximum (K_{safe}). This crucially maximizes the denominator of the residual term, tightly bounding the maximum residual disturbance to $D_{UUB}|_{\mathcal{B}_{sat}} = D_c + D_R^2 (4 m_{\min} k_{d,\min} K_{safe})^{-1}$.

For the Lyapunov derivative to remain strictly negative ($\dot{V} < 0$), the operational states must exceed the region dominated by this saturated residual boundary. Solving the quadratic tracking root yields the minimized Ultimate Bound for the tracking error (b_r):

$$\|\bar{\mathbf{r}}_c\| > \frac{D_R + \sqrt{D_R^2 + 4 m_{\min} k_{d,\min} K_{safe} D_c}}{2 m_{\min} k_{d,\min} K_{safe}} \equiv b_r \tag{28}$$

Similarly, substituting the dynamic leakage $\kappa_{max} = \hat{c}_{1,max}/Z_B$ and solving for the parameter estimation matrix at the saturation limit strictly tightens the parameter Ultimate Bound (b_Z) to its theoretical threshold:

$$\|\tilde{\mathbf{Z}}\|_F\big|_{\mathcal{B}_{sat}} > C_Z + \sqrt{\frac{D_{UUB}|_{\mathcal{B}_{sat}} Z_B}{\hat{c}_{1,max} r_{peak}}} \equiv b_Z \tag{29}$$

Evaluating b_r and b_Z mathematically guarantees that the system is Uniformly Ultimately Bounded (UUB). By dynamically maximizing the damping denominator precisely when disturbance transients peak, the composite error bounds are tightly compressed. Concurrently, the saturation of the dynamic leakage term explicitly prevents the parameter estimation matrices ($\|\tilde{\mathbf{Z}}\|_F$) from drifting. This ensures robust impact stabilization without saturating the physical actuators.

Finally, to formally extend this continuous-time stability to the complete hybrid operational domain, the overall system stability must incorporate both the continuous flow mapping and the discrete jump mapping. During the inter-impact continuous flows ($t \in (t_k, t_{k+1})$), the system energy is strictly dissipated outside the bounded region ($\dot{V} < 0$ for $\|\bar{\mathbf{r}}_c\| > b_r$). At the instantaneous discrete jumps t_k (i.e., non-differentiable rigid collisions and discontinuous state-machine phase transitions), the system inherently yields non-positive discrete energy variations ($\Delta V(t_k) = V(t_k^+) - V(t_k^-) \leq 0$), as physically established in Assumption 1 and further detailed in the subsequent Remark. Because the Lyapunov candidate function strictly decreases outside the ultimate bound during continuous flows and does not increase during discrete state jumps, the overall hybrid architecture mathematically prevents artificial energy accumulation. Thus, the theoretical UUB guarantee is rigorously preserved across all transitional landing phases and mode switching events. □

Remark 1. ***Discrete Gain Expansion at Contact Phase:*** *To aggressively reject surface impacts and enforce ZMP anchoring during the landing phase, the outer-loop stabilization gains* ($\mathbf{K}_d, \mathbf{K}_p$) *and the adaptation rate matrices* ($\mathbf{\Gamma}_W, \mathbf{\Gamma}_V$) *are discretely scaled up upon triggering the hardware contact flag. Specifically, as determined by the Optuna hyperparameter framework, the neural adaptation rate is subjected to an aggressive transient multiplier* ($\alpha_{mrac} = 28.0$)*. This forces the network to learn and suppress the rigid impact discontinuities at a severely accelerated rate before physical rollover can occur. Since this transition acts as a strict state latch, the parameters become piecewise-constant, maintaining* $\dot{\mathbf{\Gamma}} = \mathbf{0}$ *within each respective phase. Consequently, the analytical cancellation in the Lyapunov derivative* ($\dot{V}$) *remains strictly valid. Furthermore, scaling up* $\mathbf{\Gamma}$ *induces a discrete, non-positive energy jump in the Lyapunov candidate function* ($V(t_c^+) < V(t_c^-)$) *due to the* $\mathbf{\Gamma}^{-1}$ *terms in Equation (22). According to the theory of Multiple Lyapunov Functions for switched systems [23], this instantaneous energy dissipation at the switching instance guarantees that the Uniformly Ultimately Bounded (UUB) stability is strictly preserved across the entire flight-to-anchor transition without causing parameter drift.*

Remark 2. ***Actuator Limits and Sampled-Data Implementation:*** *While the theoretical proof assumes idealized unconstrained continuous control, physical actuator saturations* ($f_{\min} \leq F_i \leq f_{\max}$) *strictly restrict the maximum applicable pseudo-control effort. The dynamic bounding gain parameter* (K_{safe}) *is explicitly designed to limit the adaptive effort before these physical saturation boundaries are fundamentally breached, ensuring the control inputs remain within a practically realizable manifold. Furthermore, the discrete-time implementation executed at 1000 Hz* ($\Delta t = 1$ *ms*) *using the ETD2 scheme approximates the continuous-time dynamics with negligible discretization error, safely preserving the established UUB properties within the digital flight controller.*

3. Results

3.1. Simulation Results

The proposed adaptive architecture is evaluated using the PyBullet physics engine [24], with the Quadrotor's visual representation and structural baseline adapted from [5]. The simulation framework incorporates asymmetric spatial parameters, bidirectional actuator constraints, and tribological models detailed in Section 2.

3.1.1. Validation Setup and Environment

To explicitly isolate the contributions of the ZMP and MRAC-NN components, the validation framework evaluates structured ablations that map directly to classical linear feedback and conventional adaptive architectures. First, the **Ablation without Neural Network** functions as a classical **linear PID/PD controller**. By attempting to execute ZMP anchoring using strictly proportional-derivative feedback without adaptive learning, it isolates the failure modes of linear controllers, proving that they cannot mitigate uncompensated spatial asymmetry ($x_{com}, y_{com} \neq 0$) and ground-effect aerodynamics during impacts. Second, the **Ablation without ZMP Rules** operates precisely as a **Conventional MRAC**. By utilizing standard neural adaptation but deactivating our proposed dynamic phase-locking and bounding modifications, it proves that standard MRAC suffers from fatal parameter wind-up during non-differentiable rigid kinematic arrests.

Furthermore, optimal (e.g., **LQR**) and high-frequency robust (e.g., **SMC**) controllers are theoretically excluded from the numerical grid due to structural incompatibilities. LQR relies on Jacobian linearization around a steady-state hover, which violently diverges during the large-angle (60°) ballistic impacts evaluated in this study. Conversely, Sliding Mode Control (SMC) employs discontinuous switching functions that induce severe high-frequency chattering. When coupled with non-differentiable LCP rigid contact shocks, this chattering severely exacerbates physical bouncing and saturates actuators. Finally, the **Baseline (NonZMP)** evaluates a prevalent industry-standard heuristic (e.g., [5]) that blindly commands unconditional symmetric 100% reverse thrust without ZMP-guided geometric awareness.

To ensure deterministic contact resolution without numerical penetration, the physics engine employs a Linear Complementarity Problem (LCP) Dantzig solver [25]. The simulation time-step is fixed at $\Delta t = 1$ ms (1000 Hz) to capture high-frequency transients. The solver is configured with 2000 iterations and a split-impulse method to strictly prevent the accumulation of artificial energy during severe surface restitution.

To induce the unmodeled inertial mismatch (Δ_{dyn}) evaluated in the Lyapunov analysis, the Center of Mass (CoM) is offset from the geometric origin by randomized values of up to ± 0.12 m in the XY plane and approximately -0.03 ± 0.025 m in the Z direction. A comprehensive summary of the physical limits, neural configurations, and fused sensor noise parameters is provided in Table 1.

Practical Robustness, Optimization, and Simulation Protocol: To evaluate practical robustness against hardware constraints, we embed high-fidelity degradation models in the simulation. As detailed in Table 1, mass and inertia parameters are dynamically perturbed via truncated Gaussian distributions. Concurrently, we corrupt raw IMU and EKF states with zero-mean Gaussian noise and random-walk biases, then route them through a discrete low-pass filter to induce a deterministic 2 ms transport delay. To eliminate empirical bias under these degraded conditions, we fully automated the phase-latch logic and auxiliary parameters using the **Optuna** framework. Leveraging the CMA-ES algorithm, parameters were optimized across the 24,000-flight grid search to minimize the global penalty (J_{total}). Crucially, this revealed that applying a massive transient multiplier ($\alpha_{mrac} \approx 28.0$) to the learning rate precisely at rigid contact strictly minimizes recovery time by aggressively adapting the network before the anti-slip hinge engages. Because these severe sensor noise profiles, thrust uncertainties (e.g., asymmetric ESC degradation), and dynamic payload variations are actively injected into every single trial, the resulting 24,000-flight dataset inherently functions as a massive, multi-variable robustness evaluation rather than an idealized theoretical simulation.

Table 1. Simulation, Plant, Optimized Adaptive Control, and Sensor Noise Parameters

Parameter (Symbol)	Value	Parameter (Symbol)	Value
Plant & Actuator Constraints		**MRAC-NN Parameters**	
Nominal Mass (m_{nom})	1.0 kg (±3% dynamic unc.)	Outer Gains ($\mathbf{K}_{p_{out}}, \mathbf{K}_{d_{out}}$)	$\mathrm{diag}(0.8, 0.8, 1.15)$, $\mathrm{diag}(0.102, 0.102, 0.146)$
Inertia Unc. & Pay. Offs.	±7.5%, $\mu_{inertia} = 0.10$ m ($\sigma_{inertia} = 0.01$)	Inner Gains ($\mathbf{K}_{p_{in}}, \mathbf{K}_{d_{in}}$)	$1.7\mathbf{I}_3, 0.113\mathbf{I}_3$
Pivot Arm (l)	0.24 m	Hidden Layer Size	30 Neurons
Thrust Limits (f_{max}, f_{min})	8.3 N, −3.3 N (Rev)	Learn. Rate $\mathbf{\Gamma}_W$	$0.0014\mathrm{diag}(1, 1, 2, 10, 10, 2)$
Inertias (J_{prop}, J_{motor})	$72.6, 1040\ [\times 10^{-6}\ \mathrm{kg \cdot m^2}]$	Impact Learning Boost (α_{mrac})	×28.0 (Contact Latch)
ESC Cutoff (f_c, τ_d)	70 ± 1.4 Hz, 2 ms	Bounds (Z_B, K_{safe})	16.43, 49.30
Thrust Coeffs ($k_b, k_{b,rev}$)	$5.42, 2.06\ [\times 10^{-5}\ \mathrm{Ns^2/rad^2}]$	Leakage & Rob. (γ, σ_{c1}, K_z)	22.5, 3.0, 4.439
Torque Coeffs ($k_\tau, k_{\tau,rev}$)	$1.05, 0.63\ [\times 10^{-6}\ \mathrm{Nms^2/rad^2}]$	Impact PD Gain Mult. (α_{pd})	×1.7875
Thrust Cone Limits ($a^{\pm}_{max}, \phi_{max}$)	18, −17.7 $\mathrm{m/s^2}$, $\pi/3$ rad	Dyn. Gain Params (η_s, k_s)	32.86, 0.294
LCP Solver & Tribology		**Optimized ZMP & Sensor Noise (σ_{noise})**	
LCP Config	1 ms, 2000 iter, ERP: 0.1	IMU Accel. & Gyro	0.40 $\mathrm{m/s^2}$, 0.015 rad/s
Friction ($\mu_s, \mu_k, \mu_r, \mu_{sp}$)	(0.85, 0.6, 0.02, 0.02)/ (0.30, 0.21, 0.004, 0.002)	EKF Attitude & Pos.	0.005 rad, 0.02 m
Restitution (e_{res})	0.20/0.27 (inelastic)	EKF Rates	0.15 $\mathrm{rad/s^2}$, 0.05 m/s
Contact (k, d)	20 kN/m, 546/700 Ns/m	ZMP FHAN (r, h)	6100.0, 0.127
LuGre Stiffness (σ_{f0})	5000/2600 N/m	Barrier Gain ($\beta_{zmp}, \epsilon, \alpha_{zmp}$)	42.0, 0.02, 0.1
LuGre Damping (σ_{f1})	40.0/19.2 Ns/m	ZMP Base/Boost (G_{base}, G_{boost})	0.0145, 0.43
Viscous & Stribeck (σ_{f2}, v_s)	0.01 Ns/m, 0.03 m/s	Sat. & Damp ($K_{inv\_comp}, K_{d_{div}}$)	4.0, 0.664
Vacuum Clutch (C_{max}, C_{base})	0.96, 0.23	Kinematic & Snap (K_{kin}, C_{snap})	0.45, 0.72
Anti-Slip Rev. Idle (Ω_{anti_slip})	174.1 rad/s	Gravity Scale (K_{grav_scale})	0.75

To ensure independent reproducibility, the grid search evaluated impact velocities from 0.1 to 6.1 m/s (21 linear increments). Incline boundaries spanned 2.5° to 60° for high-friction and up to 26° for low-friction surfaces (13 steps each). Pure vertical drops were initialized parallel to the Earth frame (0°), whereas ballistic impacts are aligned with the local surface normal. A strict failure criterion (J_{fail}) was defined as platform separation or structural tilt exceeding 120° ($q_x^2 + q_y^2 > 0.75$). Successful stabilization required arresting spatial drift within a 30 s horizon using a single-layer 30-neuron MRAC-NN. Finally, to substantiate the 1000 Hz real-time feasibility claim, computational benchmarking on an AMD Ryzen 9 5900X desktop platform confirmed that the standalone $\mathcal{O}(1)$ derivative-free cascaded control loop executes in merely ≈ 0.15 ms per cycle. While the worst-case overall simulation load, including intense multi-point LCP physical contact resolution, reaches 0.45 ms, real-world flight deployment inherently bypasses the physics solver. Because the proposed Lewis MRAC formulation strictly avoids computationally expensive Jacobian inversions and iterative backpropagation, its algebraic complexity and floating-point operations (FLOPs) requirement remain exceedingly low. Consequently, the algorithm is mathematically highly suitable for deployment on standard embedded flight microcontrollers (e.g., ARM Cortex-M series architectures), possessing the computational efficiency to safely execute well within the 1 ms temporal horizon, thereby guaranteeing practical hardware feasibility for future experimental research.

3.1.2. Extensive Parameter Sweep Validation

Simulation Assumptions and Boundary Constraints: Before detailing the empirical evaluations, the following simulation boundaries are explicitly defined to contextualize the physical scope of the validation: (1) aerodynamic lateral crosswinds are assumed to be negligible during the highly transient, sub-second touchdown sequence; (2) the inclined landing platform constitutes a uniform, planar surface without irregular terrain features; and (3) structural material yielding or chassis deformations during extreme hard impacts are neglected, evaluating the airframe strictly as a rigid multibody.

To quantitatively evaluate controller performance across over 24,000 discrete impact simulations, a cumulative scalar penalty function J_{total} is computed. Unrecoverable structural rollovers receive a strict absolute penalty (J_{fail} = 20,000). Successful adherences are penalized linearly based on post-impact absolute planar drift (d, clamped to 10 m) and stabilization time (t, clamped to 30 s):

$$J_{\text{succ}} = 500\min(|d|, 10) + 100\min(\max(t, 0), 30) \tag{30}$$

The global performance metric evaluates the absolute accumulated cost across all trials in the grid without masking individual failures: $J_{\text{total}} = \sum J_i$.

The penalty scalarization coefficients and saturation limits were deterministically engineered to enforce strict operational priorities rather than arbitrary scalarization. The absolute failure penalty (20,000) is deliberately scaled to mathematically dominate the theoretical worst-case success penalty ($500 \times 10 + 100 \times 30 = 8000$). This guarantees that the evaluation metric unconditionally prioritizes structural survival over rapid settling. Within the success domain, penalizing spatial drift (500) more heavily than stabilization time (100) reflects the geometric constraints of finite landing platforms (e.g., marine vessels), where sliding off an edge poses a significantly more critical mission threat than prolonged in-place stabilization. The saturation limits (10 m and 30 s) represent these practical physical boundaries and prevent penalty gradients from diverging during extended sliding anomalies. Independent sensitivity analyses confirmed that while alternative scalar weightings shift the absolute numerical magnitude of J_{total}, the relative performance hierarchy and survivable envelopes between the proposed architecture and the ablation baselines remain strictly invariant. Furthermore, to fulfill the requirement for disaggregated performance tracking without obfuscating the global dataset, the absolute planar drift (m) and stabilization time (s) distributions are explicitly visualized as separated metrics via the auxiliary colorbars adjacent to each principal heatmap.

1. High-Friction Regimes: Vertical and Angled Impacts (Task 1 & Task 2)

On high-friction surfaces ($\mu_s = 0.85$), lateral sliding is naturally restricted, making structural rollover the dominant failure mode. To directly address the ablation requirements and quantitatively evaluate the isolated contributions of the proposed control components, both pure vertical drops (Task 1, up to 6.1 m/s) and angled ballistic trajectories (Task 2, up to 7.9 m/s) were evaluated across four distinct configurations over the entire operational grid. The configurations include the prevalent **Baseline** (symmetric reverse-thrust without ZMP), an **Ablation without Neural Network** (ZMP anchoring governed by a non-adaptive PD controller), an **Ablation without ZMP Rules** (MRAC-ZMP operating without dynamic gain expansion and phase-latch constraints), and the **Proposed Full Architecture**.

As observed in Figure 4 for pure vertical descents (Task 1), the **Baseline (NonZMP)** heuristic strategy yields a massive cumulative penalty. Its safe envelope collapses abruptly; while it survives 60° inclines at minimal speeds, the boundary plunges drastically past 4.5 m/s. Isolating the neural network, the **Ablation (No NN)** configuration relies solely on a classical PD controller attempting to track the ZMP feedforward torque. Uncompensated ground-effect aerodynamics and inertial shifts (Δ_{dyn}) overwhelm the linear controller, resulting in severe phase lag and widespread failure across steep inclines. Conversely, restoring the MRAC but deactivating the non-differentiable phase-latches and dynamic gain expansion (**Ablation-No ZMP Rules**) triggers severe parameter wind-up upon kinematic arrest. Without the anti-slip mechanical hinge, the quadrotor physically slips and rolls over, accumulating a high penalty. The **Proposed (ZMP + NN + Rules)** architecture successfully merges all subsystems, explicitly preventing wind-up while generating instantaneous asymmetric torque. This synergistic control pushes the survivable envelope to extreme 60° inclines for velocities up to 4.6 m/s, minimizing the global penalty to an unparalleled J_{total} = 3,561,785.1.

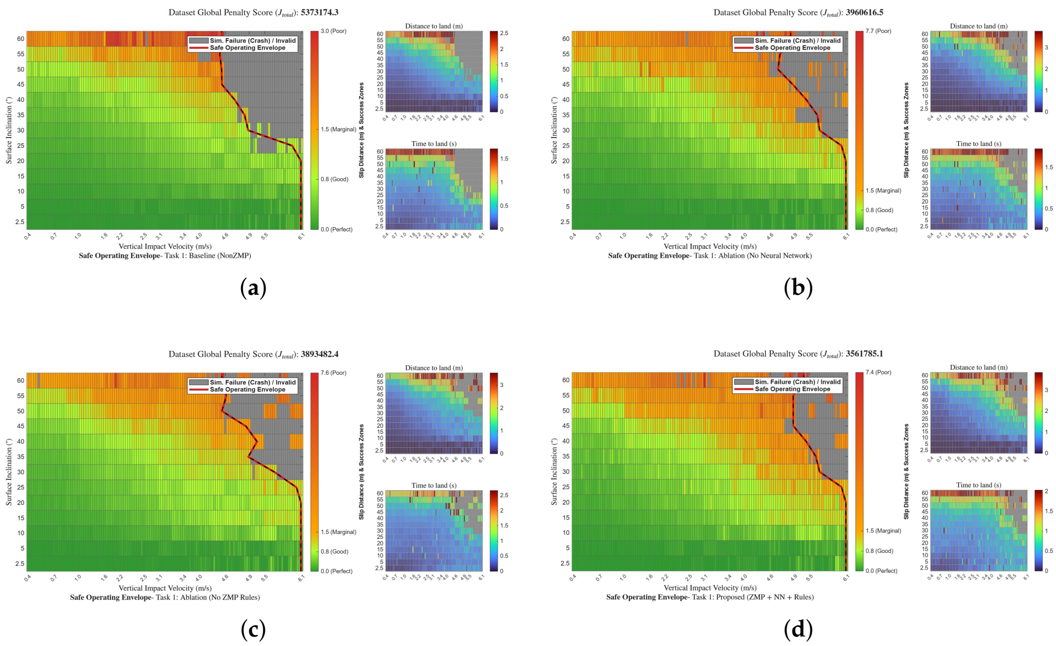

Figure 4. Task 1 (High-Friction Vertical): Cumulative penalty heatmaps across the operational grid. The proposed architecture significantly expands the survivable envelope. (**a**) Task 1: Baseline (Vertical) [J_{total} = 5,373,174.3]. (**b**) Task 1: Ablation No NN (Vertical) [J_{total} = 3,960,616.5]. (**c**) Task 1: Ablation No Rules (Vertical) [J_{total} = 3,893,482.4]. (**d**) Task 1: Proposed Architecture (Vertical) [J_{total} = 3,561,785.1] .

When subjected to severe lateral velocity components that simulate angled ballistic trajectories (Task 2, Figure 5), the contact dynamics become significantly more volatile. The instantaneous Center of Mass (CoM) shifts trigger abrupt edge-tipping moments. The **Baseline** strategy is fundamentally ill-equipped to address these lateral kinetic shocks, yielding a severe global penalty of J_{total} = 13,219,026.9, with its survivability collapsing rapidly and failing entirely beyond 1.5 m/s for inclines $>$ 55°. The **Ablation (No NN)** (J_{total} = 11,745,780.2) and **Ablation (No ZMP Rules)** (J_{total} = 9,112,147.0) configurations show measurable improvements but still suffer widespread failure across mid-to-high velocity regimes due to uncompensated inertial shifts and parameter wind-up. By contrast, the **Proposed Full Architecture** successfully resolves the chaotic rebound, securing a massive penalty reduction (J_{total} = 6,936,251.2). The active asymmetric torque allocation structurally protects the survivable ballistic envelope, maintaining stability at 55° inclines up to 3.2 m/s, and preventing rollovers at 25° even at the extreme tested boundary of 7.9 m/s.

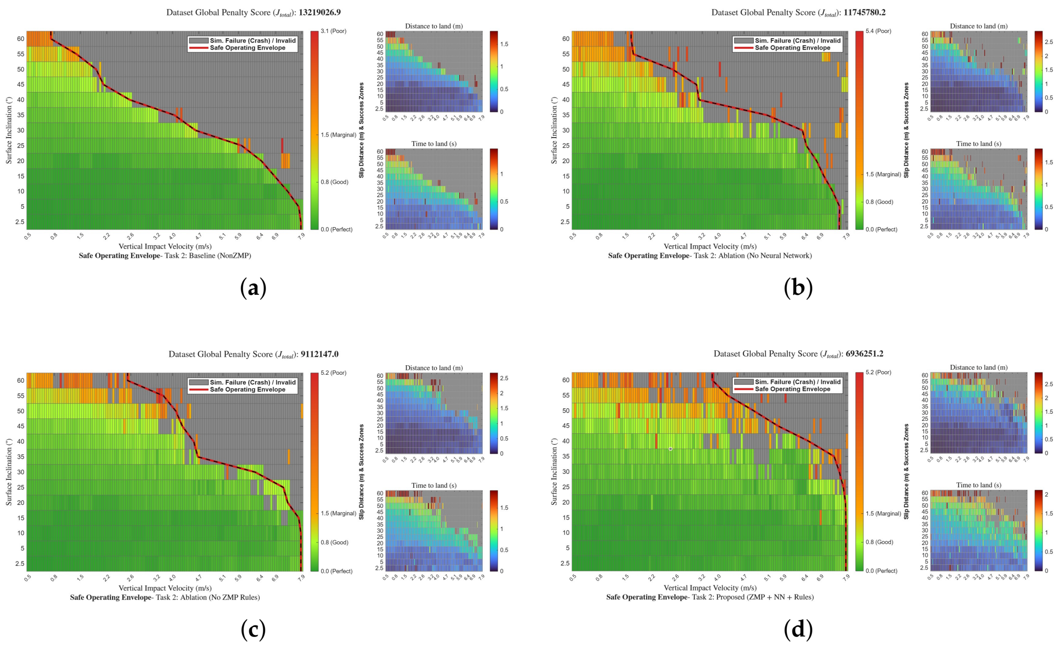

Figure 5. Task 2 (High-Friction Angled): Cumulative penalty heatmaps across the operational grid. The active asymmetric torque allocation structurally protects the survivable ballistic envelope. (**a**) Task 2: Baseline (Angled) [J_{total} = 13,219,026.9]. (**b**) Task 2: Ablation No NN (Angled) [J_{total} = 11,745,780.2]. (**c**) Task 2: Ablation No Rules (Angled) [J_{total} = 9,112,147.0]. (**d**) Task 2: Proposed Architecture (Angled) [J_{total} = 6,936,251.2].

To meet disaggregated performance-tracking requirements and validate the estimator, Tables 2 and 3 summarize the computational simulation times and ZMP tracking errors strictly for the vertical impact evaluations (Task 1).

Table 2. Computational Feasibility and CPU Simulation Time: Task 1 (High-Friction Vertical).

Control Scenario	Mean Time (s)	Std Dev (s)	Max Time (s)
Proposed (ZMP + NN)	**15.03**	**3.55**	**37.15**
Ablation (No Rules)	16.69	4.97	48.34
Ablation (No NN)	15.88	4.87	56.20
Baseline (NonZMP)	15.32	3.94	41.55

Table 3. ZMP Estimation Error Validation: Task 1. Mean Err denotes the RMSE against the analytic ground-truth. (Baseline is excluded as it relies on a non-ZMP heuristic).

Control Scenario	Mean Err (m)	Worst Err (m)
Proposed (ZMP + NN)	0.1466	**2.8348**
Ablation (No Rules)	0.1325	3.5848
Ablation (No NN)	0.1288	4.0594

Table 3 validates the bounded tracking performance of the ZMP estimator. The transient spikes in the absolute estimation error are physically inevitable, driven primarily by instantaneous LCP contact shocks at touchdown and injected high-frequency IMU measurement noise. Despite these extreme conditions, the proposed architecture bounds the maximum transient shock spike to 2.8348 m, significantly outperforming the unmitigated No-NN ablation peak of 4.0594 m. It should also be noted that outside these highly transient impact windows, the continuous ZMP tracking error remains tightly bounded at nominally low steady-state values. Mean tracking errors are reported for completeness;

since all evaluated configurations share the same ZMP estimation pipeline, the marginal differences in mean error are negligible and do not imply any advantage in estimation accuracy. Furthermore, Table 2 demonstrates algorithmic efficiency. It must be explicitly noted that the simulated reference flight trajectories (excluding the landing and stabilization phases) naturally span between 30 and 40 s. Thus, the lowest mean CPU simulation execution time (15.03 s) empirically shows that, despite integrating nonlinear neural-adaptive laws, the proposed architecture runs significantly faster than real time. Degraded ablations frequently stalled the physics engine's LCP solver due to chaotic multi-axis bouncing, proving that proactive stabilization explicitly improves computational feasibility. Finally, the execution times and ZMP error bounds detailed here serve as a representative benchmark; identical performance hierarchies were consistently recorded across all subsequent tasks, and are thus omitted hereafter for brevity.

2. Low-Friction Regimes: Vertical and Angled Impacts (Task 3 & Task 4)

On highly restrictive low-friction surfaces (e.g., icy or lubricated slopes, $\mu_s = 0.30$), the absolute theoretical static equilibrium is severely capped by the physical limits of tribology. To evaluate the controller at the absolute boundaries of physics, pure vertical drops (Task 3) and angled ballistic trajectories (Task 4) were conducted.

During pure vertical drops on low friction (Task 3, Figure 6), the proposed MRAC-ZMP architecture (J_{total} = 4,176,259.4) numerically registers a higher cumulative penalty than the Baseline strategy (J_{total} = 3,198,607.4). However, this is a theoretically sound physical trade-off rather than an algorithmic deficiency. Generating active asymmetric anchoring torque inevitably induces a coupled lateral shear force. Because static friction is strictly insufficient, the quadrotor successfully avoids structural rollover but enters a prolonged planar slide. In these extreme boundary cases, the finite geometric dimensions of the simulated landing platform are occasionally insufficient to arrest the extended drift, causing the drone to slide off the edge—registering mathematically as a catastrophic boundary failure. Furthermore, rare simulation outliers (e.g., minor kinematic clipping bugs during prolonged edge sliding) artificially inflated the cumulative penalty score. Thus, at the absolute limits of physics, the proposed method actively trades immediate catastrophic rollover for extended lateral sliding, representing a necessary and intended physical compromise.

Conversely, when subjected to lateral momentum simulating angled ballistic trajectories (Task 4, Figure 7), the structural superiority of the proposed framework becomes unequivocally evident. The Baseline strategy suffers from compounded dynamics, with its survivability envelope collapsing abruptly past 5.0 m/s, deteriorating to merely 10° at 7.7 m/s. By contrast, the MRAC-ZMP architecture capitalizes on the transient spike in normal force during the angled collision to momentarily restore fractional frictional authority. This structural advantage dramatically reduces the cumulative penalty by 32.2% (J_{total} = 3,101,938.3 vs. 4,575,391.2) and successfully prevents the high-velocity envelope collapse. The proposed method securely protects the operating envelope, maintaining stability at 22° inclines up to 5.4 m/s and sustaining an impressive 20° limit even at the extreme tested boundary of 7.7 m/s. Therefore, the MRAC-ZMP algorithm systematically extracts the maximum mathematically possible stability margin from the lubricated surface, preventing catastrophic ballistic rollovers across the critical mid-to-high velocity spectrum where the baseline consistently fails.

https://doi.org/10.3390/electronics15184130

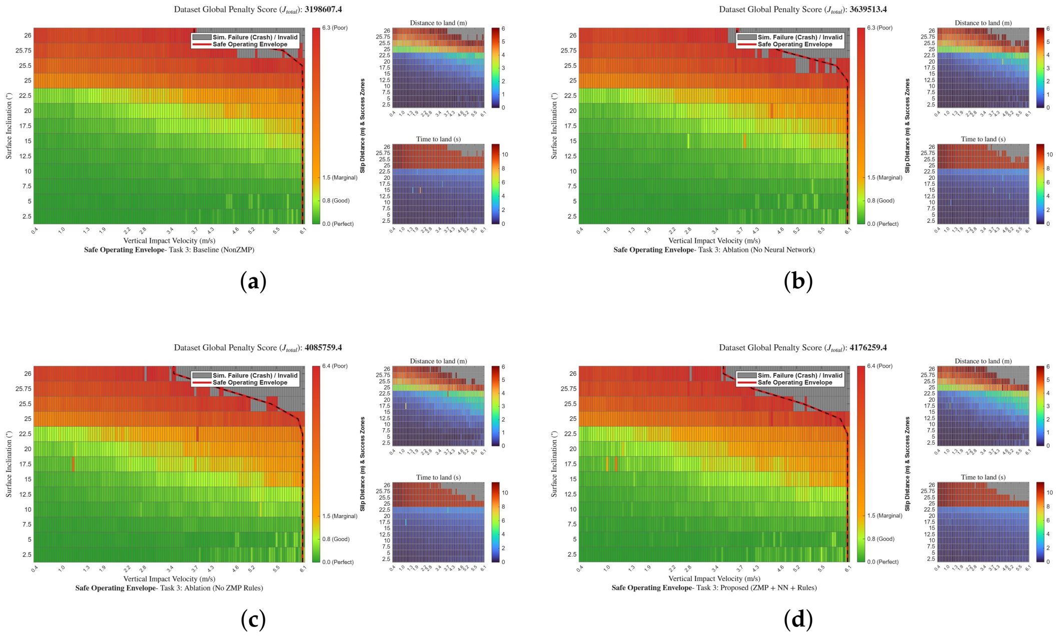

Figure 6. Task 3 (Low-Friction Vertical): Cumulative penalty heatmaps across the operational grid. The proposed method actively trades immediate catastrophic rollover for extended lateral sliding. (**a**) Task 3: Baseline (Vertical) [J_{total} = 3,198,607.4]. (**b**) Task 3: Ablation No NN (Vertical) [J_{total} = 3,639,513.4]. (**c**) Task 3: Ablation No Rules (Vertical) [J_{total} = 4,085,759.4]. (**d**) Task 3: Proposed Architecture (Vertical) [J_{total} = 4,176,259.4].

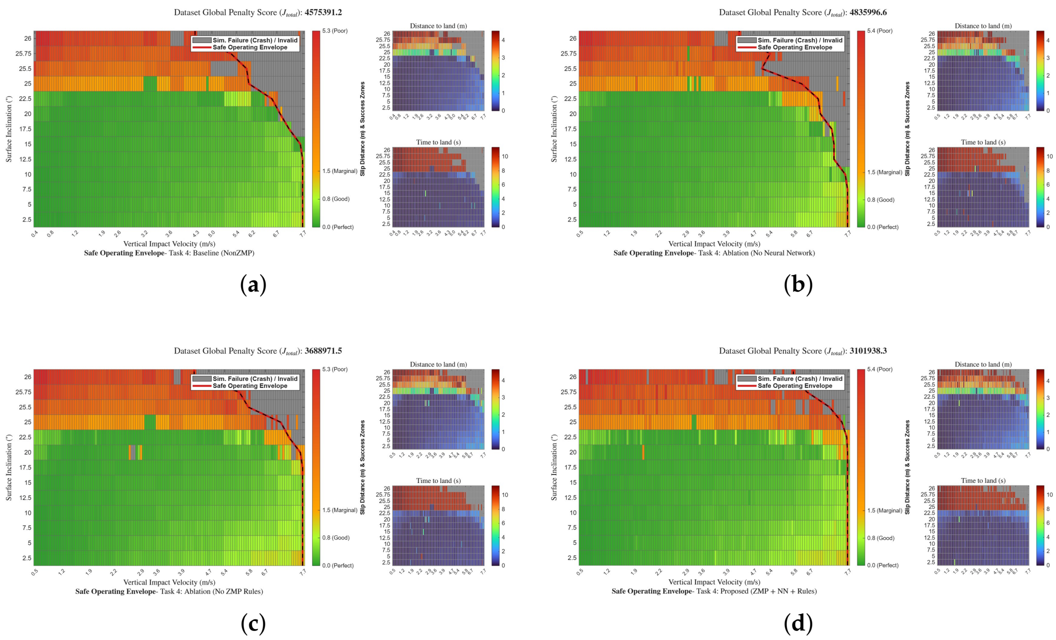

Figure 7. Task 4 (Low-Friction Angled): Cumulative penalty heatmaps across the operational grid. The MRAC-ZMP algorithm systematically extracts the maximum mathematically possible stability margin. (**a**) Task 4: Baseline (Angled) [J_{total} = 4,575,391.2]. (**b**) Task 4: Ablation No NN (Angled) [J_{total} = 4,835,996.6]. (**c**) Task 4: Ablation No Rules (Angled) [J_{total} = 3,688,971.5]. (**d**) Task 4: Proposed Architecture (Angled) [J_{total} = 3,101,938.3].

4. Discussion

The comprehensive 24,000-flight parameter sweep confirms that integrating a ZMP-guided geometric anchoring strategy with an impact-resilient MRAC-NN resolves the rollover risks inherent to inclined quadrotor landings. By framing the touchdown sequence as a real-time, ZMP-bounded adaptive control task, the system circumvents the need for computationally intensive terrain predictive models or mechanical shock absorbers.

The theoretical formulation addresses the vulnerability of standard continuous error manifolds to rigid impacts. Incorporating a kinematic reference state-locking mechanism, along with dynamically expanding bounding gains and leakage modifications, immunizes the network against parameter wind-up caused by sudden kinematic arrests. As mathematically confirmed by the Lyapunov analysis, this structural bounding guarantees Uniform Ultimate Boundedness (UUB) against severe contact shocks. Operating in parallel, the ZMP algorithm evaluates overturning moments directly in the body frame and executes immediate asymmetric torque commands, modulated by a kinetic energy compressor, to maintain the ground reaction force within the support polygon.

Validation via the LCP-Dantzig physics engine demonstrates the architecture's operational superiority over standard mitigation strategies. The grid-search results reveal that baseline heuristics relying on unconditional symmetric reverse thrust fail to dissipate asymmetric kinetic energy, resulting in broad, unrecoverable rollover zones and substantial cumulative penalties. In high-friction regimes ($\mu_s = 0.85$), the active MRAC-ZMP torque allocation converts destructive edge-tipping momentum into controlled, energy-dissipating planar slides, reducing cumulative stabilization costs by up to 47.5% and expanding the absolute survivable flight envelope to extreme 60° inclines and ballistic impact velocities up to 7.9 m/s.

Furthermore, validation in highly restrictive low-friction regimes ($\mu_s = 0.30$) highlights the system's ability to operate near the absolute physical limits of tribology. While the requisite ZMP anchoring torques inevitably induce planar sliding due to lateral shear forces on lubricated surfaces, the architecture successfully leverages transient normal-force spikes during angled collisions to maximize the theoretically achievable stability margin. In these extreme boundary cases, it is crucial to explicitly distinguish between a reduced scalar penalty and improved physical survivability. Although the proposed method occasionally incurs a higher numerical penalty due to prolonged stabilization times during extended lateral slides, it actively prevents critical ballistic rollovers where baseline controllers fail. Thus, the architecture effectively trades an immediate catastrophic structural failure for a controlled planar drift, ensuring ultimate physical survivability across the mid-to-high velocity spectrum.

Ultimately, this control architecture safely dissipates extreme kinetic energy through active planar adherence, outperforming conventional symmetric reverse-thrust methods. While the simulated physics environment natively incorporates ESC slew-rate limits and EKF transport delays, transitioning this dynamic anchoring framework to physical micro-aerial vehicle (MAV) platforms remains a primary focus for future research. Practical deployment will necessitate addressing specific hardware implementation bottlenecks, notably the phase lag introduced by onboard IMU noise filtering during violent impact shocks, and the inherent aerodynamic switching delays of bidirectional ESCs during rapid reverse-thrust zero-crossing maneuvers. Future research will focus on integrating visual-inertial odometry for preemptive surface friction estimation, extending the framework to reject sustained aerodynamic crosswinds, adapting the LCP solver to accommodate non-planar irregular terrain geometries, and adapting the proposed dynamic anchoring framework to these physical MAV constraints.

Author Contributions: Conceptualization, S.B.; methodology, Ö.A.; software, Ö.A.; validation, Ö.A.; formal analysis, Ö.A.; investigation, Ö.A.; data curation, Ö.A.; writing—original draft preparation, Ö.A.; writing—review and editing, S.B.; visualization, Ö.A.; supervision, S.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: During the preparation of this manuscript, the authors used large language models (Google Gemini 3.1 and xAI Grok-3) for the purposes of English language polishing, grammatical corrections, and LaTeX formatting. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

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

Abbreviations

The following abbreviations are used in this manuscript:

BEMT	Blade Element Momentum Theory
BET	Blade Element Theory
CLSTM-AE	Conditional Long Short-Term Memory Autoencoder
CMA-ES	Covariance Matrix Adaptation Evolution Strategy
CoM	Center of Mass
CPU	Central Processing Unit
DoF	Degrees of Freedom
EKF	Extended Kalman Filter
EMA	Exponential Moving Average
ERP	Error Reduction Parameter
ESC	Electronic Speed Controller
ETD2	Second-Order Exponentially Time-Differenced
FHAN	Fast Nonlinear Tracking Differentiator
FLOPs	Floating-Point Operations
FPU	Floating-Point Unit
GE	Ground Effect
GRF	Ground Reaction Force
IIR	Infinite Impulse Response
IMU	Inertial Measurement Unit
LCP	Linear Complementarity Problem
LPF	Low-Pass Filter
LQC	Linear Quadratic Control
LQR	Linear Quadratic Regulator
LSTM	Long Short-Term Memory
MAV	Micro-Aerial Vehicle
MPC	Model Predictive Control
MRAC	Model Reference Adaptive Control
MRAC-NN	Model Reference Adaptive Control Neural Network
NN	Neural Network
NNFL	Neural Network Feedback Linearization
OSQP	Operator Splitting Quadratic Program
PD	Proportional-Derivative
PE	Persistency of Excitation
PID	Proportional-Integral-Derivative
QP	Quadratic Programming
RMSE	Root Mean Square Error

SMC	Sliding Mode Control
TD	Tracking Differentiator
UAV	Unmanned Aerial Vehicle
UUB	Uniform Ultimate Boundedness
VRS	Vortex Ring State
ZMP	Zero Moment Point
ZOH	Zero-Order Hold

Appendix A. Kinematic Transformations and Rigid-Body Inertia

The quadrotor spatial orientation follows the Z-Y-X Euler sequence $\boldsymbol{\eta} = [\phi, \theta, \psi]^T$. The rotation matrix $\mathbf{R}_b^e \in SO(3)$ mapping the body frame ($\mathcal{F}_b$) to the Earth frame ($\mathcal{F}_e$), and the Euler rate transformation $\mathbf{T}_{eul}$ ($\boldsymbol{\omega}^b = \mathbf{T}_{eul}\dot{\boldsymbol{\eta}}$), are:

$$\mathbf{R}_b^e = \begin{bmatrix} c_\psi c_\theta & c_\psi s_\phi s_\theta - c_\phi s_\psi & s_\phi s_\psi + c_\phi c_\psi s_\theta \\ s_\psi c_\theta & c_\phi c_\psi + s_\phi s_\psi s_\theta & c_\phi s_\psi s_\theta - c_\psi s_\phi \\ -s_\theta & c_\theta s_\phi & c_\phi c_\theta \end{bmatrix} \tag{A1}$$

$$\mathbf{T}_{eul} = \begin{bmatrix} 1 & 0 & -s_\theta \\ 0 & c_\phi & c_\theta s_\phi \\ 0 & -s_\phi & c_\theta c_\phi \end{bmatrix} \tag{A2}$$

where $c_{(\cdot)}$ and $s_{(\cdot)}$ denote $\cos(\cdot)$ and $\sin(\cdot)$. The reverse rotation is $\mathbf{R}_e^b = (\mathbf{R}_b^e)^T$.

Appendix A.1. Body-Frame Kinetics and Hybrid Spatial Transformation

Rigid-body kinetics are defined relative to the geometric center of the base plane ($\mathbf{p}_{base}^b = \mathbf{0}$). The physical platform has two spatial asymmetries: Center of Mass (CoM) displacement ($\mathbf{p}_{com}^b$) and impact pivot offset ($\mathbf{p}_{piv}^b$). The base inertia tensor relates to the CoM inertia ($\mathbf{I}_{com}^b$) via the Parallel Axis Theorem: $\mathbf{I}_{base}^b = \mathbf{I}_{com}^b - m[\mathbf{p}_{com}^b]_\times^2$. Defining the coupled pseudo-inertia as $\mathbf{I}_{piv}^b = \mathbf{I}_{base}^b + m[\mathbf{p}_{piv}^b]_\times[\mathbf{p}_{com}^b]_\times$, the body-frame mass-inertia matrix ($\mathbf{M}^b$) and the velocity dependent wrench ($\mathbf{C}^b\boldsymbol{\nu}^b$) are:

$$\mathbf{M}^b = \begin{bmatrix} m\mathbf{I}_{3\times3} & -m[\mathbf{p}_{com}^b]_\times \\ m[\mathbf{p}_{com}^b - \mathbf{p}_{piv}^b]_\times & \mathbf{I}_{piv}^b \end{bmatrix} \tag{A3}$$

$$\mathbf{C}^b\boldsymbol{\nu}^b = \begin{bmatrix} m[\boldsymbol{\omega}^b]_\times\left(\mathbf{v}^b + [\boldsymbol{\omega}^b]_\times\mathbf{p}_{com}^b\right) \\ [\boldsymbol{\omega}^b]_\times\mathbf{I}_{base}^b\boldsymbol{\omega}^b + m[\mathbf{p}_{com}^b - \mathbf{p}_{piv}^b]_\times\left([\boldsymbol{\omega}^b]_\times\mathbf{v}^b\right) \\ -\, m[\mathbf{p}_{piv}^b]_\times\left([\boldsymbol{\omega}^b]_\times^2\mathbf{p}_{com}^b\right) \end{bmatrix} \tag{A4}$$

Mapping these body-fixed mechanics to the hybrid Newton–Euler state vector $\boldsymbol{\nu} = [(\mathbf{v}^e)^T, (\boldsymbol{\omega}^b)^T]^T$ defined in Section 2, the block-diagonal transformation matrix $\mathbf{T}_h \in \mathbb{R}^{6\times6}$ is applied. Mapping the hybrid twist to the body twist ($\boldsymbol{\nu}^b = \mathbf{T}_h\boldsymbol{\nu}$) and projecting the resulting wrench via $\mathbf{T}_h^T$, the transition matrices are:

$$\mathbf{T}_h = \begin{bmatrix} \mathbf{R}_e^b & \mathbf{0}_{3\times3} \\ \mathbf{0}_{3\times3} & \mathbf{I}_{3\times3} \end{bmatrix}, \quad \dot{\mathbf{T}}_h = \begin{bmatrix} -[\boldsymbol{\omega}^b]_\times\mathbf{R}_e^b & \mathbf{0}_{3\times3} \\ \mathbf{0}_{3\times3} & \mathbf{0}_{3\times3} \end{bmatrix} \tag{A5}$$

Applying the spatial congruence transformation ($\mathbf{M} = \mathbf{T}_h^T\mathbf{M}^b\mathbf{T}_h$) yields the unified hybrid mass-inertia matrix:

$$\mathbf{M} = \begin{bmatrix} m\mathbf{I}_{3\times3} & -m\mathbf{R}_b^e[\mathbf{p}_{com}^b]_\times \\ m[\mathbf{p}_{com}^b - \mathbf{p}_{piv}^b]_\times\mathbf{R}_e^b & \mathbf{I}_{piv}^b \end{bmatrix} \tag{A6}$$

Differentiating the velocity mapping ($\dot{\boldsymbol{\nu}}^b = \mathbf{T}_h \dot{\boldsymbol{\nu}} + \dot{\mathbf{T}}_h \boldsymbol{\nu}$) yields the exact hybrid velocity dependent wrench:

$$\mathbf{C}\boldsymbol{\nu} = \mathbf{T}_h^T \left(\mathbf{C}^b \mathbf{T}_h \boldsymbol{\nu} + \mathbf{M}^b \dot{\mathbf{T}}_h \boldsymbol{\nu} \right) \tag{A7}$$

Appendix A.2. Nominal Controller Vs. High-Fidelity Plant

While these equations describe the exact simulation mechanics, the algorithmic controller uses a simplified rigid-body model to ensure $\mathcal{O}(1)$ execution without numerical inversions of dense matrices. The baseline controller strictly assumes a purely diagonal inertia tensor ($\mathbf{I}^b_{base} \approx \mathrm{diag}(I_{xx}, I_{yy}, I_{zz})$) and zero spatial offsets ($\mathbf{p}^b_{com} = \mathbf{p}^b_{piv} = \mathbf{0}$). Applying these symmetric geometric assumptions to the hybrid Newton–Euler equations mathematically cancels the coupled translational-rotational Coriolis forces. This yields the block-diagonal nominal mass-inertia ($\mathbf{M}_{nom}$) and velocity-dependent ($\mathbf{C}_{nom}$) matrices utilized by the adaptive baseline:

$$\mathbf{M}_{nom} = \begin{bmatrix} m\mathbf{I}_{3\times3} & \mathbf{0}_{3\times3} \\ \mathbf{0}_{3\times3} & \mathrm{diag}(I_{xx}, I_{yy}, I_{zz}) \end{bmatrix} \tag{A8}$$

$$\mathbf{C}_{nom}\boldsymbol{\nu} = \begin{bmatrix} \mathbf{0}_{3\times1} \\ [\boldsymbol{\omega}^b]_\times \mathrm{diag}(I_{xx}, I_{yy}, I_{zz}) \boldsymbol{\omega}^b \end{bmatrix} \tag{A9}$$

Because the physical quadrotor possesses spatial asymmetries and unmodeled non-diagonal inertia elements, the resulting dynamic mismatch ($\Delta_{dyn} = (\mathbf{M}_{nom} - \mathbf{M})\dot{\boldsymbol{\nu}} + (\mathbf{C}_{nom} - \mathbf{C})\boldsymbol{\nu}$) acts as a highly coupled state-dependent disturbance. This structural disparity explicitly bounds the performance of standard cascaded controllers, mathematically necessitating the autonomous MRAC compensation framework formulated in Section 2.3.3.

Appendix B. Disturbance Wrench and Control Allocation

The hybrid disturbance wrench $\mathbf{G} \in \mathbb{R}^6$ aggregates external forces in $\mathcal{F}_e$ and moments in $\mathcal{F}_b$ ($\|\mathbf{G}\| \leq G_{max}$). It decomposes into four primary components: $\mathbf{G} = \mathbf{G}_{grav} + \mathbf{G}_{gyro} + \mathbf{G}_{aero} + \mathbf{G}_{ground}$.

Appendix B.1. Kinematic Gravity and Grounding Disturbance

Gravity acts statically in the inertial frame ($\mathbf{F}^e_{grav} = [0, 0, -mg]^T$). Since rotational dynamics are evaluated relative to the impact pivot, the spatial offset between the CoM and the pivot generates an overturning moment:

$$\mathbf{G}_{grav} = \begin{bmatrix} \mathbf{F}^e_{grav} \\ (\mathbf{p}^b_{com} - \mathbf{p}^b_{piv}) \times (\mathbf{R}^b_e \mathbf{F}^e_{grav}) \end{bmatrix} \tag{A10}$$

Similarly, proximity to the surface induces an aerodynamic grounding disturbance ($\mathbf{F}^e_{ground}$). Distinct from the rigid LCP contact constraints, this represents the asymmetric air-cushioning pressure acting on the airframe due to Ground Effect (GE), consistent with the Cheeseman–Bennett dynamics [26] formulated in Appendix B.2. As the rigid body rotates around the pivot constraint during impact, this aerodynamic grounding force maps to an equivalent geometric moment:

$$\mathbf{G}_{ground} = \begin{bmatrix} \mathbf{F}^e_{ground} \\ (\mathbf{p}^b_{com} - \mathbf{p}^b_{piv}) \times (\mathbf{R}^b_e \mathbf{F}^e_{ground}) \end{bmatrix} \tag{A11}$$

Appendix B.2. Gyroscopic Reactions and Aerodynamics

The gyroscopic wrench captures reactive moments from bidirectional rotors [27], including precession and reactive torques from rotor accelerations ($\dot{\Omega}_i$). Given propeller inertias $J_{r,i}$ and $\mathbf{e}_3 = [0,0,1]^T$ in a symmetric configuration ($s_i \in \{1,-1,1,-1\}$), this acts as a pure rotational disturbance:

$$\mathbf{G}_{gyro} = \begin{bmatrix} \mathbf{0}_{3\times1} \\ (\boldsymbol{\omega}^b \times \mathbf{e}_3)\sum_{i=1}^4 s_i J_{r,i}\Omega_i + \mathbf{e}_3 \sum_{i=1}^4 s_i J_{r,i}\dot{\Omega}_i \end{bmatrix} \tag{A12}$$

Rotor aerodynamics are modeled using Blade Element Momentum Theory (BEMT) [28] and linear-quadratic parasitic drag ($\mathbf{D}_{v_{1,2}}, \mathbf{D}_{\omega_{1,2}}$). The body-frame velocity is dynamically extracted as $\mathbf{R}_e^b \mathbf{v}^e$. The aerodynamic force ($\mathbf{F}^b_{aero}$) and torque ($\boldsymbol{\tau}^b_{aero}$) are:

$$\mathbf{F}^b_{aero} = -\big(\mathbf{D}_{v_1} + \mathbf{D}_{v_2}\mathrm{diag}(|\mathbf{R}_e^b \mathbf{v}^e|)\big)\mathbf{R}_e^b \mathbf{v}^e + \mathbf{F}^b_{flap} \tag{A13}$$

$$\boldsymbol{\tau}^b_{aero} = -\big(\mathbf{D}_{\omega_1} + \mathbf{D}_{\omega_2}\mathrm{diag}(|\boldsymbol{\omega}^b|)\big)\boldsymbol{\omega}^b + \boldsymbol{\tau}^b_{flap} \tag{A14}$$

Rotating the force component into the inertial frame completes the hybrid aerodynamic wrench: $\mathbf{G}_{aero} = [\ (\mathbf{R}_b^e \mathbf{F}^b_{aero})^T,\ (\boldsymbol{\tau}^b_{aero})^T\ ]^T$.

The BEMT module computes local inflow (λ_i) and advance (μ_i) ratios, and solves for the induced velocity ($v_{i,ind}$) via Newton–Raphson iteration of Glauert's implicit equation, which accommodates Vortex Ring State (VRS) regimes. Reactive yaw torque partitions into profile drag (k_{prof}) and induced drag: $\tau_{z_i} = [k_{prof}\Omega_i^2 + (v_{z_i} + v_{i,ind})F_i/(|\Omega_i| + \epsilon)]s_i$. Horizontal velocities induce asymmetric blade lift [29], generating parasitic flapping moments ($\mathbf{F}^b_{flap}, \boldsymbol{\tau}^b_{flap}$).

Ground Effect (GE) is incorporated via the Cheeseman–Bennett formulation [26]. A Ray-Casting algorithm projects virtual rays along $-\mathbf{R}_b^e \mathbf{e}_3$ to extract local altitude h_i, modifying thrust:

$$F_{i,true} = F_{i,cmd}\left[1 + \left(\frac{R^2}{16h_i^2 - R^2}\right)\sqrt{\frac{v_{i,ind}^2}{\|\mathbf{v}^e_{xy}\|^2 + v_{i,ind}^2}}\right] \tag{A15}$$

Appendix B.3. Asymmetric Control Allocation ($\boldsymbol{\beta}$)

The mapping of rotor thrusts to the chassis depends on the CoM location (x_{com}, y_{com}). The physical plant ($\boldsymbol{\beta}_{true}$) incorporates CoM shifts, impact pivot offsets (x_{piv}, y_{piv}), and aerodynamic ESC degradation in reverse thrust. Conversely, the nominal controller assumes a purely symmetric origin ($\boldsymbol{\beta}_{nom}$). Letting $x_{c,p} = x_{piv} - x_{com}$ and $y_{c,p} = y_{piv} - y_{com}$ denote the planar offsets, the baselines are:

$$\boldsymbol{\beta}_{true} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ -y_{c,p} & -l - y_{c,p} & -y_{c,p} & l - y_{c,p} \\ x_{c,p} - l & x_{c,p} & l + x_{c,p} & x_{c,p} \\ C_{y,1}(\cdot) & C_{y,2}(\cdot) & C_{y,3}(\cdot) & C_{y,4}(\cdot) \end{bmatrix} \tag{A16}$$

$$\boldsymbol{\beta}_{nom} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & -l & 0 & l \\ -l & 0 & l & 0 \\ -c_\tau & c_\tau & -c_\tau & c_\tau \end{bmatrix} \tag{A17}$$

In the nominal matrix, the reactive yaw torque is approximated using a static linear ratio $c_\tau = k_\tau / k_b$. However, due to the complex Blade Element Momentum Theory (BEMT) aerodynamics and asymmetric bidirectional operation, the true torque-to-thrust mapping is not a constant scalar. Instead, the transient ratio $C_{y,i}(\Omega_i, v_{i,ind}, \mathbf{v}^e)$ dynamically evolves as a highly nonlinear function of the instantaneous rotor speed, induced inflow velocities, and horizontal advance ratio. This structural mismatch ($\Delta\boldsymbol{\beta} = \boldsymbol{\beta}_{true} - \boldsymbol{\beta}_{nom}$) generates severe cross-couplings upon impact, requiring MRAC compensation. Its integration into the hybrid dynamics is mediated by the spatial projection matrix ($\mathbf{T}_h^T$).

Appendix C. Contact Jacobian and Rigid LCP Constraints

Upon contact, the quadrotor is subject to rigid constraints at the four landing gears ($\mathbf{p}^b_{leg_i}$). For an inclined surface defined by the Earth-frame normal $\mathbf{n}^e \in \mathbb{S}^2$ and orthogonal tangents $\mathbf{t}^e_1, \mathbf{t}^e_2$, localized vectors map to the body frame as $\mathbf{n}^b = \mathbf{R}^b_e \mathbf{n}^e$ and $\mathbf{t}^b_k = \mathbf{R}^b_e \mathbf{t}^e_k$.

To align constraints with the twist vector $\boldsymbol{\nu}^b$, Cartesian Jacobians are extracted using the relative contact velocity ($\mathbf{v}_{cp,i} = \mathbf{v}^b + \boldsymbol{\omega}^b \times \mathbf{p}^b_{leg_i}$). Solving the rotational cross-products, the normal ($\mathbf{J}^b_{n,i} \in \mathbb{R}^{1\times 6}$) and tangential ($\mathbf{J}^b_{t,i} \in \mathbb{R}^{2\times 6}$) Jacobians are evaluated as:

$$\mathbf{J}^b_{n,i} = \left[(\mathbf{n}^b)^T,\ (\mathbf{p}^b_{leg_i} \times \mathbf{n}^b)^T \right] \tag{A18}$$

$$\mathbf{J}^b_{t,i} = \begin{bmatrix} (\mathbf{t}^b_1)^T & (\mathbf{p}^b_{leg_i} \times \mathbf{t}^b_1)^T \\ (\mathbf{t}^b_2)^T & (\mathbf{p}^b_{leg_i} \times \mathbf{t}^b_2)^T \end{bmatrix} \tag{A19}$$

The consolidated body-frame matrix $\mathbf{J}^b_c \in \mathbb{R}^{12\times 6}$ streamlines local computation. To strictly align these physical constraints with the hybrid Newton–Euler equations (Equation (1)), this Jacobian is mapped to the hybrid twist manifold via the spatial transition matrix: $\mathbf{J}_c = \mathbf{J}^b_c \mathbf{T}_h$. Furthermore, to integrate these contact constraints directly into the generalized closed-loop error dynamics evaluated in Section 2.4.1, the hybrid Jacobian is subsequently projected into the generalized coordinate space via $\mathbf{J}_{c,x} = \mathbf{J}_c \mathbf{J}_q$, where $\mathbf{J}_q = \mathrm{blkdiag}(\mathbf{I}_{3\times 3}, \mathbf{T}_{eul})$. Spatial algebra solvers then iterate the mapped hybrid impulses ($\mathbf{J}_c^T \boldsymbol{\lambda}$) alongside LuGre tribology [12] to enforce non-penetrating dynamics.

Appendix D. ZMP Support Polygon and Voronoi Kinetics

Appendix D.1. Analytic Ground-Truth ZMP Formulation

To rigorously evaluate the estimation accuracy presented in Table 3, an absolute analytic ground-truth ZMP is computed within the simulation framework. Unlike the onboard estimator, which assumes nominal geometric symmetry and isolates variables within the IMU frame, the ground-truth formulation evaluates the full nonlinear Newton–Euler equations mapped to the exact perturbed Center of Mass (CoM).

Let the simulated CoM offset relative to the geometric origin be $\mathbf{p}_c = [c_x, c_y, c_z]^T$. The true contact-plane proper acceleration ($\mathbf{a}^O_{prop}$), incorporating gravitational compensation, is kinematically translated to the displaced CoM ($\mathbf{a}^c_{prop}$) via Euler and centripetal accelerations:

$$\mathbf{a}^c_{prop} = \mathbf{a}^O_{prop} + \dot{\boldsymbol{\omega}}^b_{gt} \times \mathbf{p}_c + \boldsymbol{\omega}^b_{gt} \times (\boldsymbol{\omega}^b_{gt} \times \mathbf{p}_c) \tag{A20}$$

where $\boldsymbol{\omega}_{gt}^{b}$ and $\dot{\boldsymbol{\omega}}_{gt}^{b}$ are the noise-free true kinematic states extracted from the physics engine. The absolute Ground Reaction Force (GRF) is then recovered by $\mathbf{F}_{grf,true}^{b} = m_{true}\mathbf{a}_{prop}^{c} - \mathbf{F}_{thrust}^{b}$, where m_{true} denotes the exact perturbed mass and $\mathbf{F}_{thrust}^{b} = [0, 0, \sum F_{rotor,i}]^{T}$.

Similarly, the total origin-referenced inertial torque ($\boldsymbol{\tau}_{O}^{b}$)—incorporating off-diagonal cross-products via the full true inertia tensor ($\mathbf{I}_{true}$) and mapped via the parallel axis (Steiner) theorem—is utilized to evaluate the exact rigid-body moment:

$$\boldsymbol{\tau}_{O}^{b} = \mathbf{I}_{true}\dot{\boldsymbol{\omega}}_{gt}^{b} + \boldsymbol{\omega}_{gt}^{b} \times (\mathbf{I}_{true}\boldsymbol{\omega}_{gt}^{b}) + \mathbf{p}_{c} \times (m_{true}\mathbf{a}_{prop}^{c}) \tag{A21}$$

By projecting this origin moment to the contact plane (where ZMP horizontal moments vanish), the exact coordinate of the uncompensated overturning threat is isolated. Thus, the analytic ground-truth ZMP evaluated against the LCP solver is given by:

$$x_{zmp}^{gt} = \frac{l\Delta F_{thrust,y}^{b} - \tau_{O,y}^{b} - lF_{grf,x,true}^{b}}{F_{grf,z,true}^{b}}, \quad y_{zmp}^{gt} = \frac{-l\Delta F_{thrust,x}^{b} + \tau_{O,x}^{b} - lF_{grf,y,true}^{b}}{F_{grf,z,true}^{b}} \tag{A22}$$

where l is the geometric arm length and ΔF_{thrust} represents the differential thrust imbalances. To ensure an exact temporal alignment during high-frequency LCP shocks, both the estimator output and this ground-truth signal are filtered through identical discrete Fast Nonlinear Tracking Differentiators (FHAN) [14] before RMSE calculation. To prevent numerical singularities during brief near-zero normal force transients (e.g., momentary free-fall rebounds), the vertical ground reaction denominator ($F_{grf,z}^{b}$) is protected by a strict lower-bound threshold of 10^{-5} N in both the estimator and the ground-truth formulations.

Appendix D.2. Voronoi Projection and $\mathcal{C}^1$ Smoothing

The spatial penetration ($\mathbf{d}_{poly}^{b}$) of the estimated ZMP relative to the support diamond ($\mathcal{S}_{diamond}$) is evaluated via an $\mathcal{O}(1)$ Voronoi projection prior to the anchoring logic. Folding coordinates into the positive quadrant ($\bar{x} = |\hat{x}_{zmp}^{b}|, \bar{y} = |\hat{y}_{zmp}^{b}|$), the scalar boundary evaluates as $s_{raw} = \bar{x} + \bar{y} - l$.

If $s_{raw} \leq -\epsilon$ ($\epsilon = 0.02$), the ZMP is stable ($\mathbf{d}_{poly}^{b} \equiv \mathbf{0}$). Exceeding the boundary activates a $\mathcal{C}^1$-continuous Huber blending [17] to prevent derivative shock waves, yielding the smoothed penetration scalar s_{eff}. Using cross-coupling $\delta = 0.5s_{eff}$ and testing $x_{proj} = \bar{x} - \delta$, $y_{proj} = \bar{y} - \delta$, the active Voronoi region isolates as:

$$\left(|d_x^b|, |d_y^b|\right) = \begin{cases} (\bar{x},\ s_{eff} - \bar{x}), & \text{if } x_{proj} < 0 \\ (s_{eff} - \bar{y},\ \bar{y}), & \text{if } y_{proj} < 0 \\ (\delta,\ \delta), & \text{otherwise} \end{cases} \tag{A23}$$

Restoring these magnitudes to their original geometric quadrants yields $\mathbf{d}_{poly}^{b}$ (Figure A1).

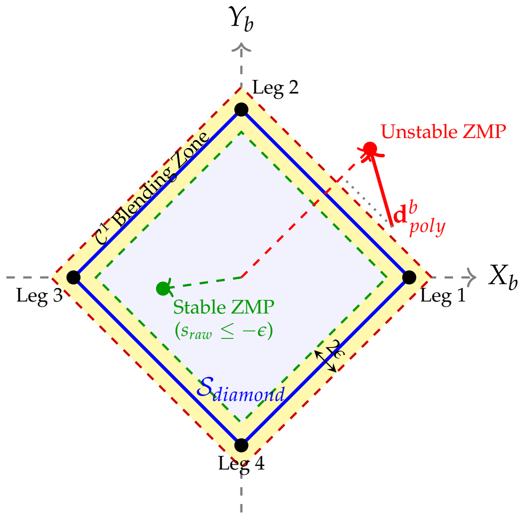

Figure A1. Quadrotor Zero Moment Point (ZMP) support polygon ($\mathcal{S}_{diamond}$) in the body frame for a '+ (Plus)' configuration. The highlighted yellow band denotes the $\mathcal{C}^1$-continuous Huber blending region ($\pm\epsilon$). The spatial penetration vector ($\mathbf{d}^b_{poly}$) analytically emerges via Voronoi projection when the measured ZMP exits this boundary, triggering feedforward torque allocation.

Appendix E. ESC and Actuator Dynamics

The framework models dynamic Electronic Speed Controllers (ESCs) and brushless motors. Commanded rotor velocities ($\Omega_{cmd,i}$) incur a transport latency (τ_d) via a ring buffer, followed by asymmetric slew-rate saturation. Emulating 'Damped Light' active braking ESCs, deceleration exploits peak available torque, while positive acceleration is limited by instantaneous Back-Electromotive Force (Back-EMF):

$$|\dot{\Omega}_i| \leq \begin{cases} \frac{\tau_{max}}{J_{r,i}}\left(1 - \frac{|\Omega_i|}{\Omega_{max_i}}\right), & \Omega_i\dot{\Omega}_i > 0 \\ \frac{\tau_{max}}{J_{r,i}}\eta_{brake}, & \Omega_i\dot{\Omega}_i < 0 \end{cases} \tag{A24}$$

where τ_{max} is the peak torque, $J_{r,i}$ is the rotor inertia, and η_{brake} is the thermal efficiency.

To replicate mechanical inertia, this rate-limited signal passes through a second-order Butterworth Infinite Impulse Response (IIR) filter. This is derived via a Zero-Order Hold (ZOH) transformation targeting a nominal cutoff frequency of $f_c = 70$ Hz.

Hardware variations are incorporated via a uniform pseudo-random tolerance array ($\pm 2\%$). This independently perturbs f_c, the maximum velocity limit (Ω_{max_i}), and $J_{r,i}$ across the actuators, generating unmodeled control allocation errors.

The filtered output is subjected to vibrations, fusing an Auto-Regressive Pink Noise model with a phase-accumulating harmonic oscillation ($\sin\left(\int_0^t \Omega_i(\tau)d\tau\right)$) that represents mechanical imbalance. The aggregated noise scales quadratically with normalized speed ($1 + 3(\Omega_i/\Omega_{max_i})^2$). By compounding latencies, Back-EMF constraints, and structural vibrations, the simulation evaluates the robust learning framework against physical actuation limits.

References

1. Wang, T.; Peng, X.; Hu, H.; Liu, T.; Xu, D. Sparse and intelligent command of unmanned swarms via hierarchical multiagent reinforcement learning. *Def. Technol.* **2026**. https://doi.org/10.1016/j.dt.2026.02.008.
2. Gräfe, A.; Eickhoff, J.; Zimmerling, M.; Trimpe, S. DMPC-Swarm: Distributed model predictive control on nano UAV swarms. *Auton. Robot.* **2025**, *49*, 28. https://doi.org/10.1007/s10514-025-10211-w.
3. Bouabdallah, S.; Murrieri, P.; Siegwart, R. Design and control of an indoor micro quadrotor. In Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Sendai, Japan, 28 September 28–2 October 2004; Volume 4, pp. 4360–4365. https://doi.org/10.1109/IROS.2004.1389659.

4. Kamel, M.; Stastny, T.; Alexis, K.; Siegwart, R. Model Predictive Control for Trajectory Tracking of Unmanned Aerial Vehicles Using Robot Operating System. In *Studies in Computational Intelligence*; Springer: Cham, Switzerland, 2017. https://doi.org/10.1007/978-3-319-54927-9_1.
5. Bass, J.; Tunney, I.; Desbiens, A. Adaptative Friction Shock Absorbers and Reverse Thrust for Fast Multirotor Landing on Inclined Surfaces. *IEEE Robot. Autom. Lett.* **2022**, *7*, 6701–6708. https://doi.org/10.1109/LRA.2022.3176102.
6. Shi, G.; Shi, X.; O'Connell, M.; Yu, R.; Azizzadenesheli, K.; Anandkumar, A.; Yue, Y.; Chung, S.J. Neural Lander: Stable Drone Landing Control Using Learned Dynamics. In Proceedings of the 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019; pp. 9784–9790. https://doi.org/10.1109/ICRA.2019.8794351.
7. González-Baldizón, Y.; Fragoso-Mandujano, J.A.; Urbina-Brito, N.; Chandomí-Castellanos, E.; Bermúdez-Rodríguez, J.I.; Pérez-Pérez, E.J.; Guzmán-Rabasa, J.A. A Hybrid Algorithm for Fault Diagnosis in Nonlinear UAV Systems Using Conditional LSTM Autoencoders. *Algorithms* **2026**, *19*, 463. https://doi.org/10.3390/a19060463.
8. Vukobratovic, M.; Borovac, B. Zero-Moment Point—Thirty Five Years of Its Life. *Int. J. Humanoid Robot.* **2004**, *1*, 157–173. https://doi.org/10.1142/S0219843604000083.
9. Qi, H.; Chen, Z.C.; Peng, Y.M.; Nie, H.; Wei, X.H.; Li, Q.Y. Research on Anti-Rollover Adaptive Landing Gear Technology Based on Impedance Control. *J. Aircr.* **2024**, *61*, 1575–1585. https://doi.org/10.2514/1.C037491.
10. Lewis, F.L.; Jagannathan, S.; Yesildirek, A. *Neural Network Control of Robot Manipulators and Non-Linear Systems*; CRC Press: Boca Raton, FL, USA, 1999; pp. 207–230.
11. Pfeiffer, F.; Glocker, C. *Multibody Dynamics with Unilateral Contacts*; John Wiley & Sons: New York, NY, USA, 1996.
12. Canudas de Wit, C.; Olsson, H.; Astrom, K.J.; Lischinsky, P. A new model for control of systems with friction. *IEEE Trans. Autom. Control* **1995**, *40*, 419–425. https://doi.org/10.1109/9.376053.
13. Lapapong, S.; Brown, A.A.; Swanson, K.S.; Brennan, S.N. Zero-moment point determination of worst-case manoeuvres leading to vehicle wheel lift. *Veh. Syst. Dyn.* **2012**, *50*, 191–214. https://doi.org/10.1080/00423114.2012.676652.
14. Han, J. From PID to Active Disturbance Rejection Control. *IEEE Trans. Ind. Electron.* **2009**, *56*, 900–906. https://doi.org/10.1109/TIE.2008.2011621.
15. Lee, T.; Leok, M.; McClamroch, N.H. Geometric tracking control of a quadrotor UAV on SE(3). In Proceedings of the 49th IEEE Conference on Decision and Control (CDC), Atlanta, GA, USA, 15–17 December 2010; pp. 5420–5425. https://doi.org/10.1109/CDC.2010.5717652.
16. Van den Bergen, G. A Fast and Robust GJK Implementation for Collision Detection of Convex Objects. *J. Graph. Tools* **1999**, *4*, 7–25. https://doi.org/10.1080/10867651.1999.10487502.
17. Boyd, S.; Vandenberghe, L. *Convex Optimization*; Cambridge University Press: Cambridge, UK, 2004.
18. Cox, S.M.; Matthews, P.C. Exponential Time Differencing for Stiff Systems. *J. Comput. Phys.* **2002**, *176*, 430–455. https://doi.org/10.1006/jcph.2002.6995.
19. Stellato, B.; Banjac, G.; Goulart, P.; Bemporad, A.; Boyd, S. OSQP: An Operator Splitting Solver for Quadratic Programs. In Proceedings of the 2018 UKACC 12th International Conference on Control (CONTROL), Sheffield, UK, 5–7 September 2018; p. 339.
20. Antonelli, G.; Cataldi, E.; Arrichiello, F.; Robuffo Giordano, P.; Chiaverini, S.; Franchi, A. Adaptive Trajectory Tracking for Quadrotor MAVs in Presence of Parameter Uncertainties and External Disturbances. *IEEE Trans. Control Syst. Technol.* **2018**, *26*, 248–254. https://doi.org/10.1109/TCST.2017.2650679.
21. Khalil, H.K. *Nonlinear Systems*, 3rd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 2002.
22. Ioannou, P.A.; Sun, J. *Robust Adaptive Control*; PTR Prentice-Hall: Upper Saddle River, NJ, USA, 1996.
23. Branicky, M.S. Multiple Lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Trans. Autom. Control* **1998**, *43*, 475–482. https://doi.org/10.1109/9.664150.
24. Coumans, E.; Bai, Y. PyBullet, a Python Module for Physics Simulation for Games, Robotics and Machine Learning. GitHub Repository. 2016. Available online: https://github.com/bulletphysics/bullet3 (accessed on 5 September 2026).
25. Stewart, D.E.; Trinkle, J.C. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and Coulomb friction. *Int. J. Numer. Methods Eng.* **1996**, *39*, 2673–2691. https://doi.org/10.1002/(SICI)1097-0207(19960815)39:15<2673::AID-NME972>3.0.CO;2-I.
26. Cheeseman, I.C.; Bennett, W.E. *The Effect of the Ground on a Helicopter Rotor in Forward Flight*; Technical Report; Aeronautical Research Council: London, UK, 1955.
27. Bresciani, T. Modelling, Identification and Control of a Quadrotor Helicopter. Master's Thesis, Lund University, Lund, Sweden, 2008.

28. Leishman, J.G. *Principles of Helicopter Aerodynamics*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2006.
29. Huang, H.; Hoffmann, G.; Waslander, S.; Tomlin, C.J. Aerodynamics and control of autonomous quadrotor helicopters in aggressive maneuvering. In Proceedings of the 2009 IEEE International Conference on Robotics and Automation, Kobe, Japan, 12–17 May 2009; pp. 3277–3282. https://doi.org/10.1109/ROBOT.2009.5152561.