Research Square

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

Coupling-Guided Nonlinear Model Predictive Control of an Aerial Manipulator under Constrained Direct Pitch Actuation

Qingling Lin
Sun Yat-sen University
Gefei Shi
shigf@mail.sysu.edu.cn

Sun Yat-sen University
Zhigang Wu
Sun Yat-sen University

Research Article

Keywords: aerial manipulator, nonlinear model predictive control, dynamic coupling, actuator allocation, constrained direct pitch actuation, base–joint coordinated motion

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11423934/v1

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

Additional Declarations: No competing interests reported.

Coupling-Guided Nonlinear Model Predictive Control of an Aerial Manipulator under Constrained Direct Pitch Actuation

Qingling Lin[a], Gefei Shi[a*], Zhigang Wu[a]

[a]School of Aeronautics and Astronautics, Sun Yat-sen University, Guangzhou, 510275, China

Abstract

This study investigates long-range end-effector target reaching for an aerial manipulator under constrained direct pitch actuation. A coupling-guided two-stage nonlinear model predictive control method is developed using a five-degree-of-freedom longitudinal prediction model derived from the coupled dynamics of a quadrotor and a two-joint manipulator. During the long-range approach, a coupling-guided allocation penalty and asymmetric actuator costs reshape the relative costs of feasible base–joint motion and input sequences, thereby reducing reliance on the direct pitch channel. Near the target, the control objective shifts to accurate end-effector regulation with base-drift suppression and prediction-terminal contraction. Comparative and ablation studies show that the proposed strategy reduces direct pitch-channel demand while maintaining comparable final positioning accuracy, and that the guidance term changes control-sequence selection rather than simply transferring pitch demand to the joints. Sensitivity analyses further demonstrate that the unloading behavior persists over the tested guidance-coefficient combinations and direct pitch-torque limits. The results indicate that base–manipulator dynamic coupling can provide an additional pathway for task realization under constrained attitude-control authority through coordinated multibody motion and actuator allocation.

Keywords: aerial manipulator; nonlinear model predictive control; dynamic coupling; actuator allocation; constrained direct pitch actuation; base–joint coordinated motion

Categories (2), (3)

*Corresponding author: Gefei Shi, email: shigf@mail.sysu.edu.cn

1. Introduction

Unmanned aerial vehicles are increasingly used beyond imaging and surveillance for inspection, maintenance, transportation, and physical interaction in complex environments[1]. Integrating robotic manipulators and task-specific sensing or actuation systems with aerial platforms has enabled increasingly complex autonomous and physical-interaction tasks[2,3]. The resulting system exhibits nonlinear, underactuated, and strongly coupled multibody dynamics[2,4]. Manipulator motion changes the mass distribution and inertia and also affects platform translation and attitude through dynamic coupling.

Existing approaches to this coupling fall broadly into modular disturbance rejection and integrated system coordination[5]. Modular schemes control the aerial platform and manipulator separately[6] and treat manipulator-induced dynamics as internal disturbances. These disturbances have been addressed through adaptive compensation of unknown coupling dynamics[7], robust control[8], adaptive backstepping[9], or coupling-model-based feedforward compensation[10]. These schemes are structurally clear, but their effectiveness can depend on the available attitude-control authority. Integrated approaches instead model the platform and manipulator as one floating-base system and explicitly retain the coupled dynamics[11]. Recent work has also emphasized predictive and constraint-aware control for aerial systems: MPC provides a systematic framework for nonlinear dynamics, input and state constraints, and finite-horizon optimization[12-14]. Predictive coordinated planning[15], sequential manipulation planning[16], and whole-body motion planning[17] further exploit the redundant degrees of freedom of the base–manipulator system. These approaches are therefore well suited to constraint-dominated aerial manipulation.

Most aerial-manipulator control strategies seek to suppress manipulator-induced reactions on the flight platform and therefore treat coupling as a disturbance[8,18]. When attitude-control authority is constrained, the available platform torque becomes an important limitation in aerial-manipulator control[19]. Continuously counteracting manipulator-induced reactions through the platform actuator may therefore place additional demand on this limited control authority. Manipulator reconfiguration and active arm motion have also been used for rotor-failure tolerance and hovering disturbance rejection through center-of-mass relocation and coupled dynamics[20,21]. These results suggest that base–manipulator coupling can provide an auxiliary dynamic regulation pathway. However, systematic use of this pathway during long-range manipulation under constrained attitude authority remains insufficiently studied.

This study considers long-range end-effector target reaching by an aerial manipulator with constrained direct pitch actuation. A coupling-guided two-stage NMPC method is proposed. Its contributions are threefold.

(1) A coupled prediction model is established for coordinated actuator allocation. The model is derived from the spatial multibody dynamics of a quadrotor with a two-joint serial manipulator. Task-plane alignment yields a five-degree-of-freedom longitudinal model that retains the essential coupling among base translation, pitch, and joint motion.

(2) A coupling-guided two-stage NMPC architecture is introduced. During the long-range approach, a stage-specific coupling-guided allocation penalty and asymmetric actuator costs change the relative costs of feasible control sequences. The optimizer consequently favors base–joint coordinated motion with lower direct pitch-torque demand. Near the target, stage switching, a base-drift penalty, and a prediction-terminal contraction condition support accurate end-effector positioning.

(3) A Fair Lock-3DOF baseline and guidance-term ablation isolate the effects of active joints and allocation guidance. Common-window statistics, guidance-coefficient sensitivity analysis, and a pitch-torque-limit scan characterize base–joint coordination and direct pitch-channel use.

2. Modeling of the Aerial Manipulation System

2.1. System description and kinematics

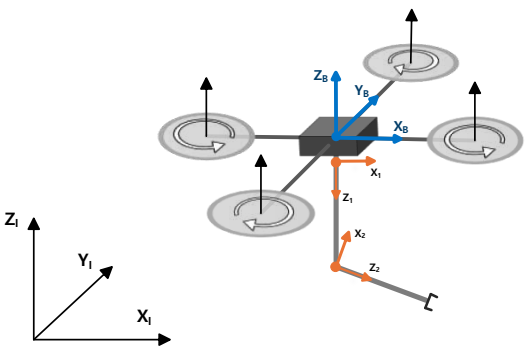


Fig. 1 Coordinate frames of the aerial manipulation system

Define the inertial frame $\{I\}=\{O_I,x_I,y_I,z_I\}$, the body-fixed frame $\{B\}=\{O_B,x_B,y_B,z_B\}$ at the platform center of mass, and the local frames $\{L_1\}=\{O_1,x_1,y_1,z_1\}$ and $\{L_2\}=\{O_2,x_2,y_2,z_2\}$ attached to the first and second links, respectively. The coordinate frames are shown in **Fig. 1**.

The platform center-of-mass position and ZYX Euler angles are, respectively,

$$\boldsymbol{p}_b^I=\left[x,y,z\right]^{\mathrm{T}} \tag{1}$$

$$\boldsymbol{\Theta}=\left[\phi,\theta,\psi\right]^{\mathrm{T}} \tag{2}$$

The generalized coordinates of the full system are defined as

$$\boldsymbol{\eta}=\left[x,y,z,\phi,\theta,\psi,q_1,q_2\right]^{\mathrm{T}}\in\mathbb{R}^{8}. \tag{3}$$

The rotation matrix from the body frame to the inertial frame is

$$\boldsymbol{R}_B^I=\boldsymbol{R}_z(\psi)\boldsymbol{R}_y(\theta)\boldsymbol{R}_x(\phi) \tag{4}$$

where $c\equiv\cos(\cdot)$, $s\equiv\sin(\cdot)$.

The manipulator mounting point relative to the platform center of mass is

$$\boldsymbol{r}_b^B=\left[0,0,-d\right]^{\mathrm{T}} \tag{5}$$

Let the two link lengths be l_1 and l_2, respectively. The serial-chain geometry yields the link center-of-mass and end-effector positions relative to the body. The inertial-frame end-effector position is

$$\boldsymbol{p}_e^I=\boldsymbol{p}_b^I+R_B^I\boldsymbol{r}_e^B \tag{6}$$

Differentiating the center-of-mass positions gives the corresponding linear velocities:

$$\boldsymbol{v}_i^I=\boldsymbol{J}_i(\boldsymbol{\eta})\dot{\boldsymbol{\eta}} \tag{7}$$

2.2. Dynamic model

The nonlinear dynamics of the complete aerial manipulation system are derived from the Euler–Lagrange equation with generalized viscous damping:

$$\frac{d}{dt}\left(\frac{\partial L}{\partial\dot{\boldsymbol{\eta}}}\right)-\frac{\partial L}{\partial\boldsymbol{\eta}}+\boldsymbol{D}\dot{\boldsymbol{\eta}}=\boldsymbol{Q} \tag{8}$$

$$L=T-U \tag{9}$$

The total kinetic energy T and its components are

$$\begin{aligned}T&=T_b+\sum_{i=1}^{2}T_i\\ T_b&=\frac{1}{2}m_b\boldsymbol{v}_b^{I^T}\boldsymbol{v}_b^I+\frac{1}{2}\boldsymbol{\omega}_b{}^T\boldsymbol{I}_b\boldsymbol{\omega}_b\\ T_i&=\frac{1}{2}m_i\boldsymbol{v}_i^{I^T}\boldsymbol{v}_i^I+\frac{1}{2}\boldsymbol{\omega}_i{}^T\boldsymbol{I}_i\boldsymbol{\omega}_i\end{aligned} \tag{10}$$

Here, T, T_b, and T_i are the kinetic energies of the complete system, quadrotor, and each manipulator link. m_b and $\boldsymbol{I}_b$ denote the mass and inertia tensor of the quadrotor platform, respectively, whereas m_i and $\boldsymbol{I}_i$ denote the mass and inertia tensor of the i th manipulator link. The total potential energy U and its components are

$$\begin{aligned}U&=U_b+\sum_{i=1}^{2}U_i\\ U_b&=m_b g\boldsymbol{e}_z^T\boldsymbol{p}_b^I\\ U_i&=m_i g\boldsymbol{e}_z^T(\boldsymbol{p}_b^I+\boldsymbol{R}_b^I\boldsymbol{r}_i^B)\end{aligned} \tag{11}$$

Here, U, U_b, and U_i are the potential energies of the complete system, quadrotor, and each link, and g is gravitational acceleration. The Lagrange–d'Alembert equation gives the full dynamics as

$$\boldsymbol{M}(\boldsymbol{\eta})\ddot{\boldsymbol{\eta}}+\boldsymbol{h}(\boldsymbol{\eta},\dot{\boldsymbol{\eta}})+\boldsymbol{D}\dot{\boldsymbol{\eta}}=\boldsymbol{Q} \tag{12}$$

$$\boldsymbol{M}_{ij}(\boldsymbol{\eta}) = \frac{\partial^2 L}{\partial \dot{\boldsymbol{\eta}}_i \partial \dot{\boldsymbol{\eta}}_j} \tag{13}$$

$$\boldsymbol{h}(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}}) = \frac{\partial}{\partial \boldsymbol{\eta}} \left(\frac{\partial T}{\partial \dot{\boldsymbol{\eta}}} \right) \dot{\boldsymbol{\eta}} - \frac{\partial T}{\partial \boldsymbol{\eta}} + \frac{\partial U}{\partial \boldsymbol{\eta}} \tag{14}$$

Substituting the generalized coordinate vector $\boldsymbol{\eta} = [x, y, z, \phi, \theta, \psi, q_1, q_2]^{\mathrm{T}}$ gives the mass matrix $\boldsymbol{M}(\boldsymbol{\eta})$, the combined Coriolis, centrifugal, configuration-dependent inertial, and gravitational term $\boldsymbol{h}(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}})$, and the generalized viscous damping matrix $\boldsymbol{D}^{8\times 8}$. The total thrust T acts along the body $+z_B$ direction, the differential-thrust body torque is $\boldsymbol{\tau}_B = [\tau_\phi, \tau_\theta, \tau_\psi]^{\mathrm{T}}$, and the generalized force vector $\boldsymbol{Q}$ is

$$\boldsymbol{Q} = \begin{bmatrix} T\boldsymbol{R}_B^I \boldsymbol{e}_z \\ \boldsymbol{E}^T(\phi, \theta)\boldsymbol{\tau}_B \\ \tau_1 \\ \tau_2 \end{bmatrix} \tag{15}$$

The joint torques act on generalized coordinates q_1 and q_2. Because the generalized coordinates use the absolute platform pose and relative joint angles, the virtual work of the joint actuators enters the joint coordinates directly. Its effect on platform attitude appears through the coupling terms in $\boldsymbol{M}(\boldsymbol{\eta})$ and $\boldsymbol{h}(\boldsymbol{\eta}, \dot{\boldsymbol{\eta}})$.

2.3. Task-plane alignment and longitudinal dynamics

To retain the dominant base–manipulator coupling while reducing the NMPC problem dimension, the platform is yaw-aligned with the target before coordinated control. Let the horizontal target displacement from the initial platform position be $\Delta x = x_d - x_0$ and $\Delta y = y_d - y_0$. The horizontal distance is

$$\rho_d = \sqrt{\Delta x^2 + \Delta y^2} \tag{16}$$

The target yaw angle is $\psi_d = \operatorname{atan2}(\Delta y, \Delta x)$. After yaw alignment, lateral, roll, and residual yaw motion outside the task plane are neglected. The five retained generalized coordinates are $\boldsymbol{q} = [x, z, \theta, q_1, q_2]^{\mathrm{T}} \in \mathbb{R}^5$. Projecting the virtual work onto this reduced coordinate set then yields the longitudinal dynamics

$$\boldsymbol{M}_{2d}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{h}_{2d}(\boldsymbol{q}, \dot{\boldsymbol{q}}) + \boldsymbol{D}_{2d}\dot{\boldsymbol{q}} = \boldsymbol{Q}_{2d} \tag{17}$$

Here, $\boldsymbol{Q}_{2d} = \left[T\sin\theta, T\cos\theta, \tau_\theta, \tau_1, \tau_2\right]^T \in \mathbb{R}^5$ is the reduced generalized-force vector, $\dot{\boldsymbol{q}}$ denotes generalized velocity, $\boldsymbol{x} = \left[\boldsymbol{q}^T, \dot{\boldsymbol{q}}^T\right]^T = \left[x, z, \theta, q_1, q_2, \dot{x}, \dot{z}, \dot{\theta}, \dot{q}_1, \dot{q}_2\right]^T \in \mathbb{R}^{10}$ is the state, and $\boldsymbol{u} = \left[T, \tau_\theta, \tau_1, \tau_2\right]^T \in \mathbb{R}^4$ is the control input. The continuous-time state-space model is

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}) = \begin{bmatrix} \dot{\boldsymbol{q}} \\ \boldsymbol{M}_{2d}^{-1}(\boldsymbol{q})\left[\boldsymbol{Q}_{2d} - \boldsymbol{h}_{2d}(\boldsymbol{q}, \dot{\boldsymbol{q}}) - \boldsymbol{D}_{2d}\dot{\boldsymbol{q}}\right] \end{bmatrix}. \tag{18}$$

Define the generalized coordinates of the base and manipulator joints as

$$\boldsymbol{q}_b = [x, z, \theta]^{\mathrm{T}}, \quad \boldsymbol{q}_m = [q_1, q_2]^{\mathrm{T}} \tag{19}$$

The five-degree-of-freedom longitudinal dynamics can then be written using the same nonlinear and damping definitions:

$$\begin{bmatrix} \boldsymbol{M}_{2d,bb} & \boldsymbol{M}_{2d,bm} \\ \boldsymbol{M}_{2d,mb} & \boldsymbol{M}_{2d,mm} \end{bmatrix} \begin{bmatrix} \ddot{\boldsymbol{q}}_b \\ \ddot{\boldsymbol{q}}_m \end{bmatrix} + \begin{bmatrix} \boldsymbol{h}_{2d,b} \\ \boldsymbol{h}_{2d,m} \end{bmatrix} + \begin{bmatrix} \boldsymbol{D}_{2d,b}\dot{\boldsymbol{q}}_b \\ \boldsymbol{D}_{2d,m}\dot{\boldsymbol{q}}_m \end{bmatrix} = \begin{bmatrix} \boldsymbol{Q}_{2d,b} \\ \boldsymbol{Q}_{2d,m} \end{bmatrix} \tag{20}$$

The generalized pitch equation is obtained as

$$\tau_\theta = M_{\theta x}\ddot{x} + M_{\theta z}\ddot{z} + M_{\theta\theta}\ddot{\theta} + M_{\theta 1}\ddot{q}_1 + M_{\theta 2}\ddot{q}_2 + h_\theta + (\boldsymbol{D}\dot{\boldsymbol{q}})_\theta \tag{21}$$

The active manipulator degrees of freedom allow the same end-effector task to be realized through different base–joint motion combinations. Configuration, velocity, and joint acceleration change the inertial and nonlinear terms in the pitch dynamics and therefore the required direct pitch torque. Joint torque does not directly or algebraically replace τ_θ. Instead, joint motion changes task realization through the complete multibody dynamics. This property motivates the coupling-guided actuator allocation.

3. Control Strategy

3.1. Two-stage constrained NMPC formulation

A fourth-order Runge–Kutta scheme discretizes the prediction model and gives the state-transition equation

$$\boldsymbol{x}_{i+1|k} = \boldsymbol{F}(\boldsymbol{x}_{i|k}, \boldsymbol{u}_{i|k}) \tag{22}$$

At each sampling instant k within the prediction horizon N, the following finite-horizon nonlinear optimal-control problem is solved from the current state $\boldsymbol{x}(t_k)$:

$$\begin{aligned} &\min_{\boldsymbol{u}_0:N-1, \boldsymbol{s}_0:N-1} J_N \\ s.t.\ \ &\boldsymbol{x}_{0|k} = \boldsymbol{x}(t_k) \\ &\boldsymbol{x}_{i+1|k} = \boldsymbol{F}(\boldsymbol{x}_{i|k}, \boldsymbol{u}_{i|k}), i = 0,1,\ldots,N-1 \\ &\boldsymbol{g}(\boldsymbol{x}_{i|k}, \boldsymbol{u}_{i|k}, \boldsymbol{s}_{i|k}) \le 0, \boldsymbol{s}_{i|k} \ge 0 \end{aligned} \tag{23}$$

Here, $g(\cdot)$ denotes the feasible set defined by input, state, and end-effector safety constraints. A single-shooting formulation is used, with the control sequence and slack variables as the main decision variables.

The stage indicator $\delta \in \{0,1\}$ is determined by the end-effector position error:

$$e_e = \left\| \boldsymbol{p}_e - \boldsymbol{p}_{e,d} \right\|_2, \qquad \delta = \begin{cases} 0, & e_e > r_{\text{near}}, \\ 1, & e_e \le r_{\text{near}}. \end{cases} \tag{24}$$

The objective function is written in the unified form

$$J_N(k) = \sum_{i=1}^{N} \left[(1-\delta)\ell_1(i) + \delta\ell_2(i) + \ell_s(i) \right] + (1-\delta)V_{f,1} + \delta V_{f,2} + l_{s,f} \tag{25}$$

Here, ℓ_1 and ℓ_2 are the Stage 1 and Stage 2 running costs; $V_{f,1}$ and $V_{f,2}$ are the stage-dependent terminal costs; and ℓ_s is the path-slack penalty.

The inputs satisfy magnitude and input-increment constraints:

$$\begin{aligned} \boldsymbol{u}_{\min} &\le \boldsymbol{u}_{i|k} \le \boldsymbol{u}_{\max} \\ \Delta\boldsymbol{u}_{\min} &\le \Delta\boldsymbol{u}_{i|k} \le \Delta\boldsymbol{u}_{\max} \end{aligned} \tag{26}$$

Nonnegative slack variables $\boldsymbol{s}_i = [s_{p,i}, s_{a,i}, s_{v,i}, s_{q,i}]^T$ soften selected position, attitude, velocity, and joint-velocity constraints. Their path penalty is

$$\ell_s = \rho_p s_p^2 + \rho_a s_a^2 + \rho_v s_v^2 + \rho_q s_q^2 \tag{27}$$

At each sampling instant, only the first element of the accepted control sequence is applied. The prediction problem is then solved again from the updated state. The overall two-stage NMPC architecture is shown in **Fig. 2**.

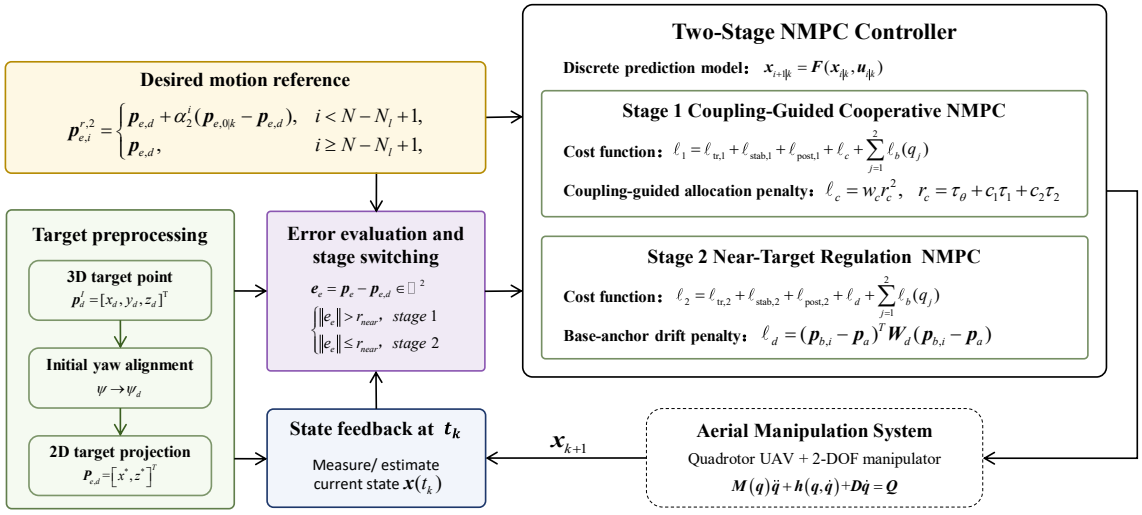


Fig. 2 Two-stage NMPC architecture

3.2. Coupling-guided cooperative approach in Stage 1

When $\|\boldsymbol{e}_e\| > r_{near}$, $r_{near} = 1.0\,\mathrm{m}$, equivalently $\delta = 0$, the control strategy operates in Stage 1. It does not require exact target arrival within a single finite prediction horizon. Instead, it generates a smooth, constrained long-range approach while using the coupled base–manipulator dynamics to shape actuator allocation.

Let the end-effector position at the start of the prediction be $\boldsymbol{p}_{e,0|k}$ and the final target be $\boldsymbol{p}_{e,d}$. To avoid an excessively aggressive transient caused by using the distant final target at every prediction node, Stage 1 uses an exponentially varying progressive reference:

$$\boldsymbol{p}_{e,i}^{r,1} = \boldsymbol{p}_{e,d} + \alpha_1^i \left(\boldsymbol{p}_{e,0|k} - \boldsymbol{p}_{e,d} \right) \tag{28}$$

Here, $0 < \alpha_1 < 1$ and $\alpha_1 = 0.72$. The Stage 1 end-effector position error is

$$\boldsymbol{e}_{e,i}^{1} = \boldsymbol{p}_{e,i} - \boldsymbol{p}_{e,i}^{r,1}, \tag{29}$$

As the prediction index i increases, $\alpha_1^i \to 0$, so the prediction reference moves progressively from the current end-effector position toward the final target.

The predicted end-effector velocity is $\boldsymbol{v}_{e,i} = \boldsymbol{J}_e(\boldsymbol{q}_i)\dot{\boldsymbol{q}}_i$. Define the hover reference input $\boldsymbol{u}_h = [m_{\mathrm{tot}} g, 0, 0, 0]^{\mathrm{T}}$ and the input deviation $\tilde{\boldsymbol{u}}_i = \boldsymbol{u}_i - \boldsymbol{u}_h$. The Stage 1 base running cost includes end-effector position and velocity errors, joint velocity, input deviation, and input increments:

$$\ell_{\mathrm{tr},1} = (\boldsymbol{e}_{e,i}^{1})^T \boldsymbol{W}_{e,1} \boldsymbol{e}_{e,i}^{1} + \dot{\boldsymbol{q}}_i^T \boldsymbol{W}_{v,1} \dot{\boldsymbol{q}}_i + \boldsymbol{v}_{e,i}^T W_{ve,1} \boldsymbol{v}_{e,i} + \tilde{\boldsymbol{u}}_i^T \boldsymbol{W}_{u,1} \tilde{\boldsymbol{u}}_i + \Delta \boldsymbol{u}_i^T \boldsymbol{W}_{\Delta u,1} \Delta \boldsymbol{u}_i \tag{30}$$

To regulate use of the constrained direct pitch channel, define the coupling-guidance residual as

$$r_{c,i} = \tau_{\theta,i} + c_1 \tau_{1,i} + c_2 \tau_{2,i} \tag{31}$$

The coupling-guided allocation penalty is then

$$\ell_{c,i} = w_c r_{c,i}^2 \tag{32}$$

The physical base–manipulator coupling is already represented completely by $\boldsymbol{M}_{2d}(\boldsymbol{q})$ and $\boldsymbol{h}_{2d}(\boldsymbol{q}, \dot{\boldsymbol{q}})$ in the prediction model. The coupling-guidance residual and its penalty provide allocation-guidance regularization; neither represents an equivalent physical coupling torque or an exact physical coupling coefficient. They change the relative costs of feasible sequences and favor base–joint coordinated motion with lower direct pitch-torque demand.

Stage 1 also penalizes attitude and altitude deviations to prevent excessive platform motion while exploiting manipulator coupling:

$$\ell_{\mathrm{stab},1} = w_{\theta,1} \theta_i^2 + w_{z,1} (z_i - z_{r,1})^2 \tag{33}$$

Here, $z_{r,1}$ is the base altitude at the beginning of the current prediction horizon. It limits unnecessary altitude changes within a single receding horizon during the long-range stage. The Stage 1 configuration cost is

$$\ell_{\mathrm{post},1} = (\boldsymbol{q}_i - \boldsymbol{q}_{r,1})^{\mathrm{T}} \boldsymbol{W}_{q,1} (\boldsymbol{q}_i - \boldsymbol{q}_{r,1}) \tag{34}$$

A joint-limit barrier discourages prolonged operation near the physical joint bounds:

$$\ell_b(q_j) = k_b \left[\varphi_\beta(q_j - q_{j,\max} + \Delta q_m) + \varphi_\beta(q_{j,\min} + \Delta q_m - q_j) \right] \tag{35}$$

Here, $\varphi_\beta(\cdot)$ is the smooth softplus function.

The complete Stage 1 running cost is therefore

$$\ell_1 = \ell_{\mathrm{tr},1} + \ell_{\mathrm{stab},1} + \ell_{\mathrm{post},1} + \ell_c + \sum_{j=1}^{2} \ell_b(q_j) \tag{36}$$

3.3. Near-target regulation and prediction-terminal contraction

When $\|\boldsymbol{e}_e\| \le r_{near}$, equivalently $\delta = 1$, the control strategy switches to Stage 2. The platform position in the longitudinal plane at switching, $\boldsymbol{p}_a = [x_s, z_s]^{\mathrm{T}}$, is stored as the base anchor. The Stage 2 reference trajectory is

$$\boldsymbol{p}_{e,i}^{r,2} = \begin{cases} \boldsymbol{p}_{e,d} + \alpha_2^i (\boldsymbol{p}_{e,0|k} - \boldsymbol{p}_{e,d}), & i < N - N_l + 1, \\ \boldsymbol{p}_{e,d}, & i \ge N - N_l + 1, \end{cases} \tag{37}$$

Here, $\alpha_2 = 0.63$. The reference still approaches the final target smoothly over the early prediction nodes. For the last 10 nodes, the reference is set directly to the final target. The Stage 2 base running cost is

$$\ell_{\mathrm{tr},2} = (\boldsymbol{e}_{e,i}^2)^T \boldsymbol{W}_{e,2} \boldsymbol{e}_{e,i}^2 + \dot{\boldsymbol{q}}_i^T W_{v,2} \dot{\boldsymbol{q}}_i + \boldsymbol{v}_{e,i}^T W_{ve,2} \boldsymbol{v}_{e,i} + \tilde{\boldsymbol{u}}_i^T \boldsymbol{W}_{u,2} \tilde{\boldsymbol{u}}_i + \Delta \boldsymbol{u}_i^T \boldsymbol{W}_{\Delta u,2} \Delta \boldsymbol{u}_i \tag{38}$$

Compared with Stage 1, Stage 2 assigns substantially larger weights to end-effector position and velocity regulation and removes the coupling-guided allocation penalty. The optimization therefore shifts from unloading the direct pitch channel to final positioning and suppression of residual motion.

The Stage 2 attitude and altitude regulation term is

$$\ell_{\mathrm{stab},2} = w_{\theta,2} \theta_i^2 + w_{z,2} (z_i - z_s)^2 \tag{39}$$

The Stage 2 configuration cost is

$$\ell_{\mathrm{post},2} = (\boldsymbol{q}_i - \boldsymbol{q}_{r,2})^{\mathrm{T}} \boldsymbol{W}_{q,2} (\boldsymbol{q}_i - \boldsymbol{q}_{r,2}) \tag{40}$$

A base-anchor drift penalty limits unnecessary platform displacement during final end-effector regulation:

$$\ell_{d,i} = (\boldsymbol{p}_{b,i} - \boldsymbol{p}_a)^T \boldsymbol{W}_d (\boldsymbol{p}_{b,i} - \boldsymbol{p}_a) \tag{41}$$

Here, $\boldsymbol{p}_{b,i} = [x_i, z_i]^T$ is a soft performance measure rather than a hard constraint.

The complete Stage 2 running cost is therefore

$$\ell_2 = \ell_{\mathrm{tr},2} + \ell_{\mathrm{stab},2} + \ell_{\mathrm{post},2} + \ell_d + \sum_{j=1}^{2} \ell_b(q_j) \tag{42}$$

A stage-dependent terminal cost is added at the end of the prediction horizon. Define

$$\boldsymbol{e}_{N|k} = \boldsymbol{p}_e(\boldsymbol{x}_{N|k}) - \boldsymbol{p}_{e,d} \tag{43}$$

$$\boldsymbol{v}_{N|k} = \boldsymbol{J}_e(\boldsymbol{q}_N) \dot{\boldsymbol{q}}_N \tag{44}$$

The terminal cost is

$$V_{f,j} = \boldsymbol{e}_{N|k}^T \boldsymbol{P}_{e,j} \boldsymbol{e}_{N|k} + \dot{\boldsymbol{q}}_N^T \boldsymbol{P}_{v,j} \dot{\boldsymbol{q}}_N + \boldsymbol{v}_{N|k}^{\mathrm{T}} \boldsymbol{P}_{ve,j} \boldsymbol{v}_{N|k}, \quad j = 1,2 \tag{45}$$

The terminal condition is not activated immediately after the switch to Stage 2. The actual end-effector error must first satisfy

$$e_e(k) \le r_f \tag{46}$$

Here, $r_f = 0.50\,\mathrm{m}$. Define a terminal radius that depends on the current closed-loop task error:

$$\varepsilon_f(k) = \max\left\{\varepsilon_{\min}, \gamma_f e_e(k)\right\}, \qquad 0 < \gamma_f < 1 \tag{47}$$

The terminal contraction factor is $\gamma_f = 0.75$. The terminal condition is

$$\left\| \boldsymbol{e}_{N|k} \right\|_2^2 \le \varepsilon_f^2(k) + s_f, \qquad s_f(k) \ge 0 \tag{48}$$

Here, $s_f(k)$ is the terminal slack variable. The terminal-slack penalty is $\ell_{s,f} = \rho_f s_f^2$, with $\rho_f = 10^6$, which imposes a strong penalty on terminal-constraint relaxation. For larger errors, the predicted terminal region contracts proportionally to the current error. As the error decreases, the radius is lower-bounded by a value close to $\varepsilon_{\min} = 0.01\,\mathrm{m}$, which reduces the risk of finite-horizon infeasibility immediately after stage switching.

3.4. Terminal error regulation and closed-loop implementation

The minimum terminal-error radius is $\varepsilon_{\min} = 0.01\,\mathrm{m}$. Far from the target, the terminal region contracts in proportion to the current task error. For sufficiently small errors, it remains inside the finite target neighborhood defined by $\varepsilon_{\min}$.

Let the predicted terminal end-effector error obtained at sampling instant k be

$$e_{N|k} = \left\| \boldsymbol{p}_e(\boldsymbol{x}_{N|k}) - \boldsymbol{p}_{e,d} \right\|_2 \tag{49}$$

When $s_f(k) = 0$,

$$e_{N|k}^2 \le \varepsilon_f^2(k) \tag{50}$$

Because both sides are nonnegative,

$$e_{N|k} \le \varepsilon_f(k) = max\{\varepsilon_{min}, \gamma_f e_e(k)\} \tag{51}$$

If the current closed-loop error also satisfies $e_e(k) > \frac{\varepsilon_{\min}}{\gamma_f}$, then

$$\gamma_f e_e(k) > \varepsilon_{min} \tag{52}$$

It follows that

$$e_{N|k} \le \gamma_f e_e(k) < e_e(k) \tag{53}$$

Assume that the terminal condition is active, the terminal slack is zero, and the current error lies outside the local region defined by $\varepsilon_{\min} / \gamma_f$. The predicted terminal state must then lie in a strictly smaller end-effector error region.

When $e_e(k) \le \frac{\varepsilon_{\min}}{\gamma_f}$,

$$e_{N|k} \le \varepsilon_{\min} \tag{54}$$

Suppose that the accepted finite-horizon candidate at sampling instant k has a finite objective value $J_N(k)$ and that a finite constant $\bar{J}$ satisfies

$$J_N(k) \le J \tag{55}$$

All objective terms are nonnegative, and the objective contains the terminal-slack penalty $\rho_f s_f^2(k)$. Therefore,

$$\rho_f s_f^2(k) \le J_N(k) \le \bar{J} \tag{56}$$

Hence,

$$0 \le s_f(k) \le \sqrt{\frac{\bar{J}}{\rho_f}} \tag{57}$$

which gives

$$e_{N|k} \le \sqrt{\varepsilon_f^2(k) + \sqrt{\frac{\bar{J}}{\rho_f}}} \tag{58}$$

With zero slack, the predicted terminal state is subject to strict contraction; nonzero slack instead yields an explicit prediction-terminal error bound. Receding-horizon control applies only the first input $\boldsymbol{u}(k) = \boldsymbol{u}_{0|k}$ and then rebuilds the finite-horizon problem from the next closed-loop state:

$$\boldsymbol{x}(k+1) = \boldsymbol{x}_{1|k} \tag{59}$$

These relations establish strict prediction-terminal contraction when terminal slack is zero and an explicit prediction-terminal error bound when slack is nonzero. They do not establish one-step closed-loop error contraction because receding-horizon control applies only the first predicted input before the problem is solved again. The resulting closed-loop positioning behavior is evaluated numerically in the next section.

4. Results and Discussion

4.1. Simulation setup and comparison methods

The initial three-dimensional position of the aerial platform is $\boldsymbol{p}_{b,0}^{I} = [0, 0, 0.5]^T$ m. The three-dimensional end-effector target is $\boldsymbol{p}_{e,d}^{I} = [10.6066, 10.6066, 3.6]^T$ m. After Stage-0 yaw alignment, the target is mapped into the longitudinal task plane, giving $\boldsymbol{p}_{e,d} = [15, 3.6]^{\mathrm{T}}$ m. The initial manipulator configuration is $q_1(0) = 10^{\circ}$ and $q_2(0) = -20^{\circ}$.

The closed-loop sampling period and prediction-horizon length are $T_s = 0.06\,\mathrm{s}$ and $N = 25$, respectively, and the total simulation time is approximately 16 s. Under nominal conditions, the direct pitch torque satisfies $|\tau_\theta| \le 0.20\,\mathrm{N\cdot m}$.

Three control strategy configurations were used to separate the effects of active manipulator degrees of freedom and coupling guidance. Coop-Guided is the proposed control strategy. It uses the full five-degree-of-freedom coupled model in both stages and includes the coupling-guided allocation penalty in Stage 1. Fair Lock-3DOF fixes the manipulator at its initial configuration during Stage 1 while retaining the manipulator mass, inertia, and configuration-dependent mass distribution in

the platform dynamics. It restores the same complete five-degree-of-freedom control strategy in Stage 2 and uses the same parameters, weights, and constraints as Coop-Guided. Coop-NoGuide uses the same prediction dynamics, constraints, stage-switching logic, and actuator costs as Coop-Guided but removes the Stage 1 guidance penalty. The Coop-Guided versus Coop-NoGuide comparison therefore isolates how the guidance term changes control-sequence selection and direct pitch-channel use.

In addition to the final end-effector error $e_{e,f} = e_e(T_{\mathrm{sim}})$, the Stage 1 pitch-torque root-mean-square value quantifies demand on the direct pitch actuator:

$$\tau_{\theta,\mathrm{RMS}}^{(1)} = \sqrt{\frac{1}{T_1}\int_0^{T_1} \tau_\theta^2(t)\,\mathrm{d}t} \tag{60}$$

The accumulated absolute pitch torque and the normalized pitch-actuator utilization in Stage 1 are

$$J_\theta^{(1)} = \int_0^{T_1} \left|\tau_\theta(t)\right| \mathrm{d}t \tag{61}$$

$$U_{\theta,RMS}^{(1)} = \frac{\tau_{\theta,\mathrm{RMS}}^{(1)}}{\tau_{\theta,\max}} \tag{62}$$

For the guidance-coefficient sensitivity study, the Stage 1 pitch-torque saturation ratio was also evaluated:

$$\rho_{\mathrm{sat}} = \frac{N_{\mathrm{sat}}}{N_1} \tag{63}$$

Here, T_1 is the Stage 1 duration, N_{sat} is the number of Stage 1 samples near the pitch-torque saturation boundary, and N_1 is the total number of Stage 1 samples. The final error $e_{e,f}$ measures task completion, whereas $\tau_{\theta,\mathrm{RMS}}^{(1)}$, $J_\theta^{(1)}$, and $U_{\theta,RMS}^{(1)}$ quantify direct pitch-channel demand during task execution.

The principal system and closed-loop simulation parameters are listed in Table 1. The NMPC parameters, weighting matrices, actuator limits, state constraints, and slack penalties used in the simulations are summarized in Table 2. Unless otherwise stated, the same physical constraints and controller parameters were used for all comparison cases.

Table 1 Principal system and closed-loop simulation parameters

Parameter	Value	Parameter	Value	Parameter	Value
m_b	1.20 kg	m_1	0.25 kg	m_2	0.20 kg
I_b	0.030 kg·m^2	I_1	2.55×10^{-3} kg·m^2	I_2	1.04×10^{-3} kg·m^2
l_1	0.35 m	l_2	0.25 m	d	0.12 m

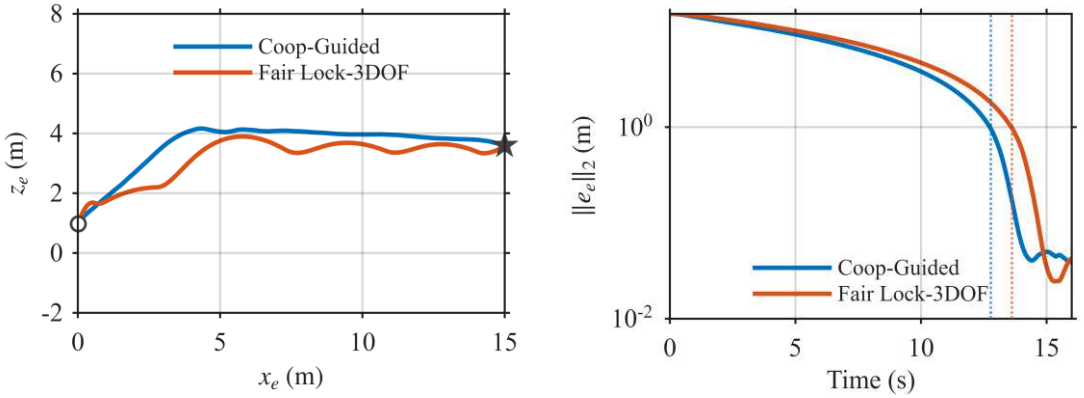
Table 2 Summary of NMPC parameters, weighting matrices, and constraints used in the simulations

Parameter	Value	Parameter	Value
w_c	1300	$\boldsymbol{W}_d$	$\mathrm{diag}(260,320)$
$\boldsymbol{W}_{e,1}$	$\mathrm{diag}(270,260)$	$\boldsymbol{W}_{e,2}$	$\mathrm{diag}(4700,3200)$
$\boldsymbol{W}_{v,1}$	$\mathrm{diag}(1.2,1.4,100,0.35,0.35)$	$\boldsymbol{W}_{v,2}$	$\mathrm{diag}(2,2,3,1.4,1.4)$
$\boldsymbol{W}_{ve,1}$	$\mathrm{diag}(18,22)$	$\boldsymbol{W}_{ve,2}$	$\mathrm{diag}(260,300)$
$\boldsymbol{W}_{u,1}$	$\mathrm{diag}(0.010,700,0.001,0.001)$	$\boldsymbol{W}_{u,2}$	$\mathrm{diag}(0.012,0.030,0.030,0.030)$
$\boldsymbol{W}_{\Delta u,1}$	$\mathrm{diag}(0.10,420,0.010,0.010)$	$\boldsymbol{W}_{\Delta u,2}$	$\mathrm{diag}(0.20,0.22,0.16,0.16)$
$\omega_{\theta,1}$, $\omega_{\theta,2}$	12000, 180	$\omega_{z,1}$, $\omega_{z,2}$	40, 25
$\boldsymbol{W}_{q,1}$	$\mathrm{diag}(0,0,100,6,6)$	$\boldsymbol{W}_{q,2}$	$\mathrm{diag}(0,0,0,1.5,1.5)$
$\boldsymbol{q}_{r,1}$	$[0,0,0,-18^\circ,18^\circ]^T$	$\boldsymbol{q}_{r,2}$	$[0,0,0,-8^\circ,8^\circ]^T$
$\boldsymbol{P}_{e,1}$	$\mathrm{diag}(185,220)$	$\boldsymbol{P}_{e,2}$	$4\mathrm{diag}(3200,4200)$
$\boldsymbol{P}_{v,1}$	$\mathrm{diag}(15,15,18,10,10)$	$\boldsymbol{P}_{v,2}$	$2.5\mathrm{diag}(120,120,180,140,140)$
$\boldsymbol{P}_{ve,1}$	$\mathrm{diag}(12,16)$	$\boldsymbol{P}_{ve,2}$	$2.5\mathrm{diag}(260,340)$
ρ_p, ρ_a, ρ_v, ρ_q	$8.0\times10^4, 1.5\times10^5, 2.0\times10^4, 500$	k_b, β	250, 60
Δq_m	8°	$\boldsymbol{D}_{2d}$	$\mathrm{diag}(0.20,0.25,0.03,0.03,0.03)$
$T_{\min}$, $T_{\max}$	$0.60m_{tot}g, 2.50m_{tot}g$	$\lvert\tau_1\rvert_{\max}$, $\lvert\tau_2\rvert_{\max}$	0.80 N·m, 0.60 N·m

$\lvert\Delta T\rvert_{\max}$, $\lvert\Delta\tau_\theta\rvert_{\max}$	3.0 N/ sample , 0.10 N·m/ sample	$\lvert\Delta\tau_1\rvert_{\max}$, $\lvert\Delta\tau_2\rvert_{\max}$	0.15 N·m/ sample
$\lvert\dot{x}\rvert_{\max}$, $\lvert\dot{z}\rvert_{\max}$	1.0 m/ s	$\lvert\theta\rvert_{\max}$, $\lvert\dot{\theta}\rvert_{\max}$	20° , 60° / s
$q_{1,2,\min/\max}$, $\lvert\dot{q}_1\rvert_{\max}$, $\lvert\dot{q}_2\rvert_{\max}$	[−80°,80°], 60° / s , 60° / s	$z_{\min}$, $p_{e,z,\min}$	0.3 m, 0.20 m

4.2. Nominal comparison and coupling-guidance ablation

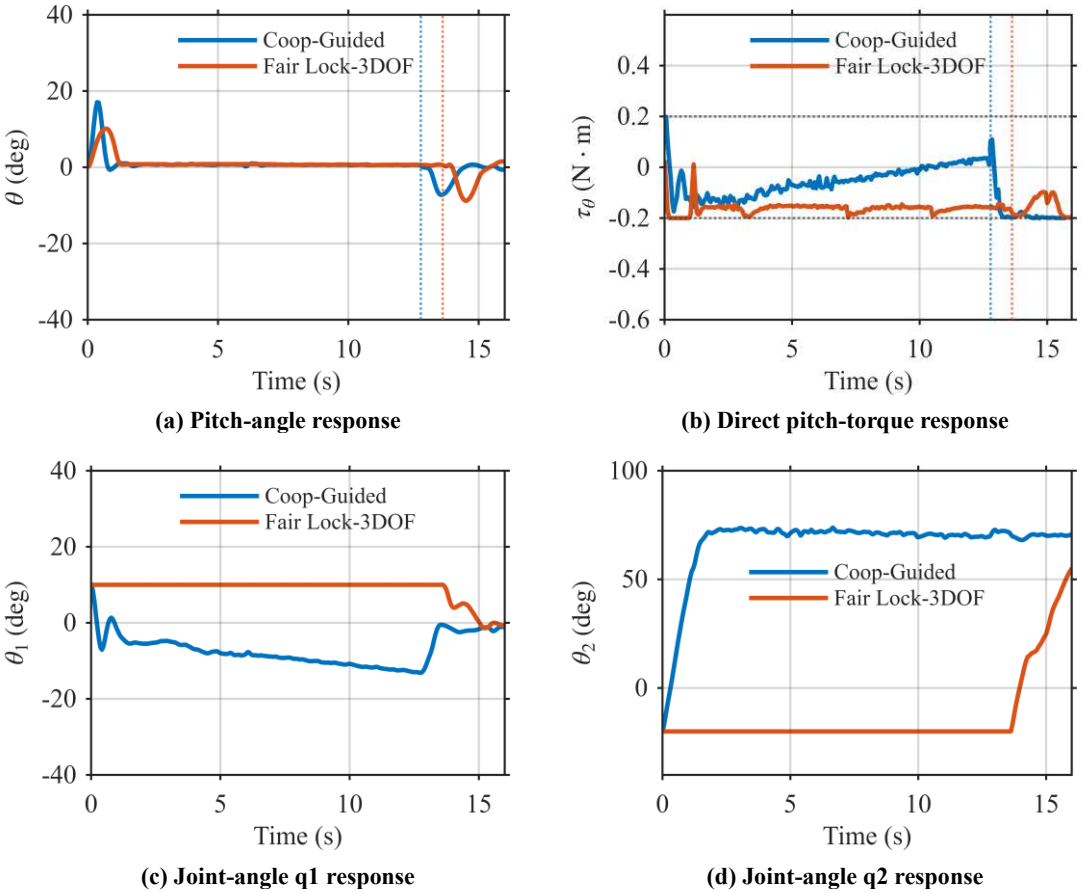
Coop-Guided and Fair Lock-3DOF were first compared under the nominal pitch-torque limit $\tau_\theta^{\max} = 0.20\,\mathrm{N\cdot m}$.



(a) End-effector trajectory **(b) End-effector error**

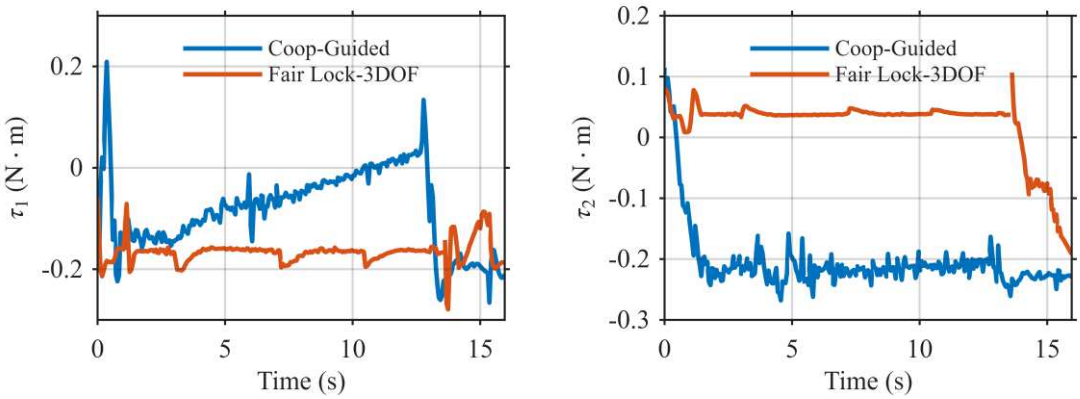
Fig. 3 End-effector trajectory and error under nominal conditions

Fig. 3 shows that Coop-Guided and Fair Lock-3DOF both entered the target region within the prescribed time. Their final end-effector errors were 0.0395 m and 0.0431 m, respectively. Accordingly, active joint participation did not compromise final positioning for the tested task.



(a) Pitch-angle response **(b) Direct pitch-torque response**

(c) Joint-angle q1 response **(d) Joint-angle q2 response**



(e) Joint-torque τ1 response **(f) Joint-torque τ2 response**

Fig. 4 Platform pitch response and manipulator joint motion under nominal conditions

Fig. 4 presents the platform pitch angle and direct pitch torque. Fair Lock-3DOF remained near the negative pitch-torque limit for much of Stage 1. With the manipulator locked, horizontal acceleration and altitude support relied primarily on platform pitch, and the associated pitch regulation was supplied directly by τ_θ. Coop-Guided remained farther from sustained saturation while performing the same task. Its Stage 1 pitch-torque RMS decreased from 0.1628 to 0.0797 N·m, a reduction of 51.1%. Accumulated absolute pitch torque decreased from 2.1925 to 0.8285 N·m·s, a reduction of 62.2%. The principal change was therefore allocation of control effort among the available actuation channels under the same pitch constraint, rather than a simple change in end-effector path.

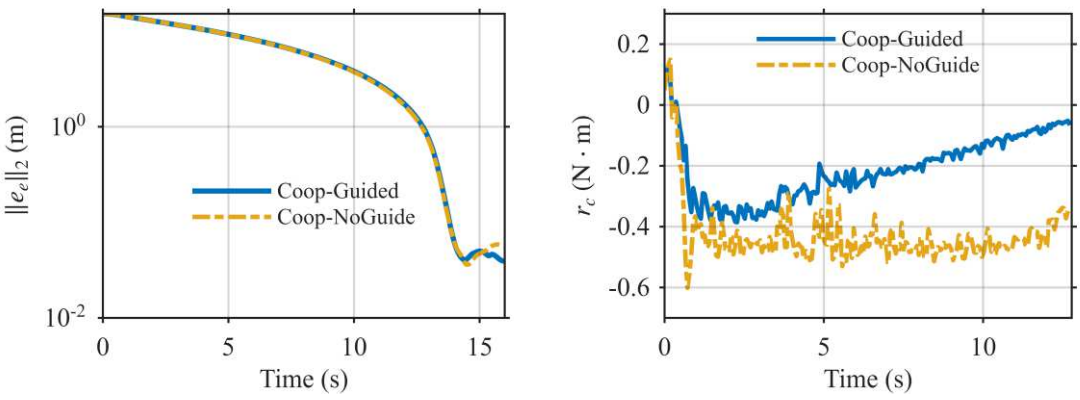
The joint responses in **Fig. 4** show continuous configuration changes and nonzero joint torques throughout Stage 1. In the Euler–Lagrange model, these motions enter the platform pitch dynamics through the configuration-dependent mass matrix $\boldsymbol{M}_{2d}(\boldsymbol{q})$ and nonlinear term $\boldsymbol{h}_{2d}(\boldsymbol{q},\dot{\boldsymbol{q}})$. Joint torque does not directly or algebraically replace τ_θ. Instead, the resulting joint motion changes joint acceleration, mass distribution, and coupled inertial response, allowing the same end-effector demand to be realized through different multibody motions. Fair Lock-3DOF fixes q_1 and q_2, leaving the optimizer to rely mainly on base translation and pitch. Coop-Guided instead coordinates base and joint motion so that the same end-effector task is realized with a different multibody motion pattern.

Coop-Guided did not minimize every attitude metric. Its maximum absolute pitch angle was approximately 17.09°, compared with 10.14° for Fair Lock-3DOF, while the full-horizon pitch-angle RMS values were 2.74° and 2.62°, respectively. The pitch-torque reduction therefore did not result from keeping the base more stationary. Instead, the control strategy permitted bounded platform attitude motion and active joint motion to expand the available multibody motion combinations. The supported conclusion is direct pitch-channel unloading through changed base–joint motion and input allocation, not simultaneous reduction of all attitude measures.

To distinguish the effect of movable joints from explicit use of coupling, Coop-Guided was compared with Coop-NoGuide. Their only difference was the coupling-guided allocation penalty:

$$\ell_c = w_c r_c^2 \tag{64}$$

$$r_c = \tau_\theta + c_1\tau_1 + c_2\tau_2 \tag{65}$$



(a) End-effector error **(b) Coupling-guidance residual**

Fig. 5 Ablation comparison between Coop-Guided and Coop-NoGuide

Fig. 5 compares the end-effector error and coupling-guidance residual. Both methods completed the target-reaching task, with final errors of 0.0395 m for Coop-Guided and 0.0565 m for Coop-NoGuide. They entered Stage 2 at 12.78 and 12.72 s, respectively, indicating that guidance did not primarily accelerate the long-range approach. Removing the guidance penalty increased the Stage 1 pitch-torque RMS from 0.0797 to 0.1810 N·m. This value was 56.0% lower with Coop-Guided. The accumulated absolute pitch torque decreased from 2.2759 to 0.8285 N·m·s, a reduction of 63.6%, and normalized utilization decreased from 0.9050 to 0.3985. The Stage 1 guidance-residual RMS also decreased from 0.4408 to 0.2347 N·m, or 46.8%.

The ablation shows that physical coupling alone does not favor control sequences with lower direct pitch-channel demand. Coop-NoGuide has the same active joints and coupled mass matrix as Coop-Guided, but without allocation guidance the optimizer assigns greater demand to τ_θ. Adding the coupling-guided allocation penalty does not change the feasible set or physical dynamics. It changes the relative costs of feasible sequences in the objective function and favors solutions with coordinated joint and base action. The corresponding nominal and ablation metrics are summarized in Table 3.

Table 3 Key quantitative metrics for the nominal comparison and ablation study.

Control strategy	e_e **(m)**	$\tau^{(1)}_{\theta,\mathrm{RMS}}$ **(N·m)**	$J^{(1)}_\theta$ **(N·m·s)**	$U^{(1)}_{\theta,RMS}$	**Stage 2 switching time (s)**
Coop-Guided	0.0395	0.0797	0.8285	0.3985	12.78
Fair Lock-3DOF	0.0431	0.1628	2.1925	0.8141	13.62
Coop-NoGuide	0.0565	0.1810	2.2759	0.9050	12.72

To remove the influence of different Stage 1 durations on accumulated metrics, the three actuator torques were evaluated over the common window $0 \le t \le 12$ s. Fair Lock-3DOF uses nonzero inverse-dynamics holding torques to maintain its locked Stage 1 joint configuration.

Table 4 Actuator-torque comparison over the common time window.

Control strategy	$\tau_{\theta,\mathrm{RMS}}$ (N·m)	$\tau_{1,\mathrm{RMS}}$ (N·m)	$\tau_{2,\mathrm{RMS}}$ (N·m)	J_θ (N·m·s)	J_1 (N·m·s)	J_2 (N·m·s)
Coop-Guided	0.0817	0.0915	0.2082	0.8057	0.8997	2.4461
Fair Lock-3DOF	0.1636	0.1691	0.0405	1.9479	2.0226	0.4739
Coop-NoGuide	0.1831	0.1963	0.2191	2.1837	2.2811	2.5784

The common-window statistics in Table 4 follow the same trend as the separate Stage 1 statistics, so the reduced direct pitch torque cannot be attributed solely to different stage durations. Relative to Fair Lock-3DOF, Coop-Guided reduced the first-joint torque RMS but increased the second-joint torque RMS. Activating the joint degrees of freedom therefore redistributed, rather than uniformly increased, joint loading. Relative to Coop-NoGuide, Coop-Guided reduced both the RMS and accumulated absolute torque in all three torque channels. The guidance term therefore did not simply transfer reduced pitch torque to larger joint torques. It changed the coordinated motion and control sequence selected for the same end-effector task.

Together, the nominal comparison, ablation, and common-window statistics support a two-part interpretation. The complete coupled dynamics provide alternative feasible motion combinations, and the guidance penalty changes their relative optimization costs. Combined, these two mechanisms reduced direct pitch-channel demand for the tested task.

4.3. Sensitivity to guidance coefficients and direct pitch-torque limits

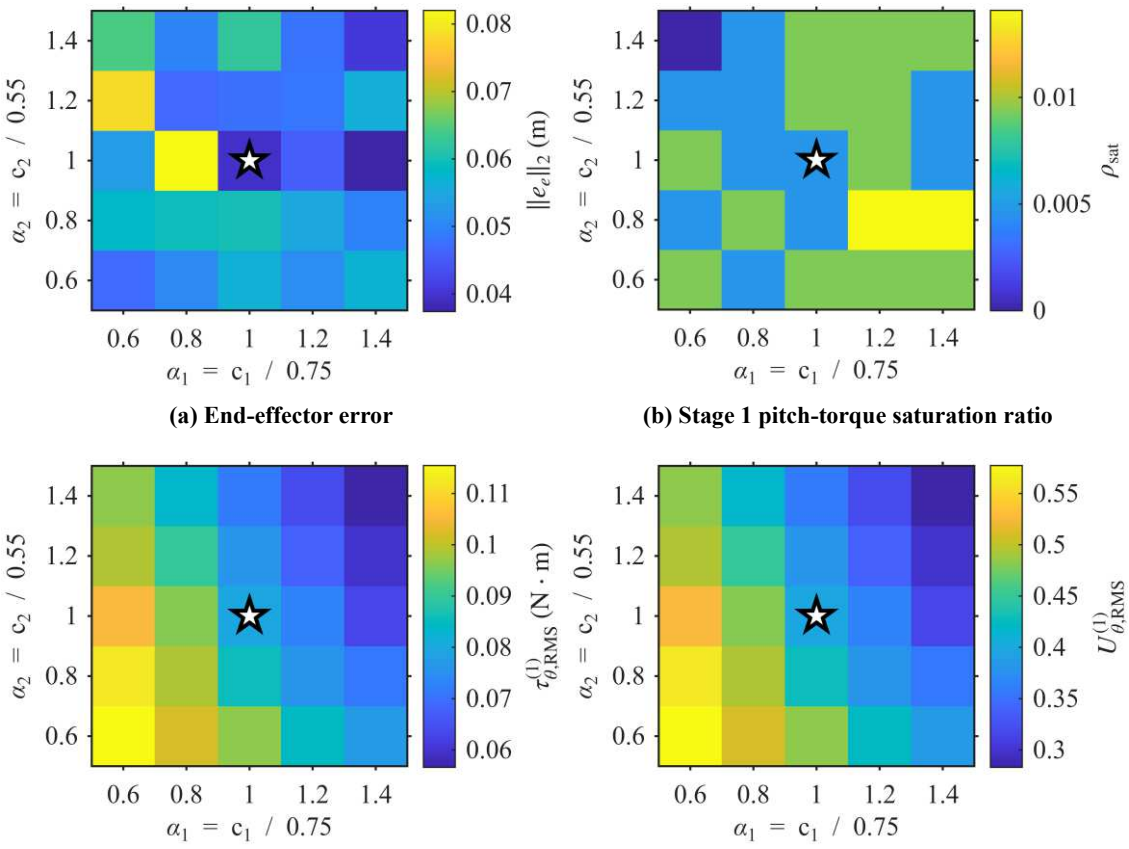
The nominal study used fixed guidance coefficients $c_1 = 0.75$ and $c_2 = 0.55$. To determine whether pitch-channel unloading depended on this specific pair, a two-dimensional sensitivity analysis varied c_1 and c_2. Define the normalized coefficients as

$$\alpha_1 = \frac{c_1}{0.75}, \quad \alpha_2 = \frac{c_2}{0.55} \tag{66}$$

Using $\alpha_1, \alpha_2 \in \{0.6, 0.8, 1.0, 1.2, 1.4\}$ produced 25 fixed-guidance coefficient combinations. The sensitivity ranges are summarized in Table 5.

Table 5 Key quantitative metrics of the guidance-coefficient sensitivity study.

Control strategy	e_e **(m)**	$\tau^{(1)}_{\theta,\mathrm{RMS}}$ **(N·m)**	$U^{(1)}_{\theta,RMS}$	ρ_{sat}
Coop-NoGuide	0.0565	0.1810	0.9050	41.98%
Tested guidance-coefficient range	0.0374–0.0820	0.0566–0.1155	0.283–0.577	0–1.41%



(a) End-effector error **(b) Stage 1 pitch-torque saturation ratio**

(c) Stage 1 pitch-torque RMS **(d) Stage 1 normalized pitch-actuator utilization**

Fig. 6 Two-dimensional sensitivity analysis of the coupling-guidance coefficients

Fig. 6 maps final error, Stage 1 saturation ratio, Stage 1 pitch-torque RMS, and normalized pitch utilization over (α_1, α_2). All 25 combinations completed the task, with final errors of 0.0374 – 0.0820 m. Basic task feasibility therefore did not require one precisely tuned coefficient pair. The guidance coefficients nevertheless affected direct pitch-channel demand. Stage 1 pitch-torque RMS ranged from 0.0566 to 0.1155 N·m, below the Coop-NoGuide value of 0.1810 N·m for every pair. Normalized utilization ranged from 0.283 to 0.577, compared with 0.9050 for Coop-NoGuide. Pitch unloading therefore persisted beyond the nominal $(\alpha_1, \alpha_2) = (1,1)$.

The heat maps show that increasing c_1 and c_2 generally reduced $\tau^{(1)}_{\theta,\mathrm{RMS}}$ and $U^{(1)}_{\theta}$. This trend is consistent with the objective-level guidance mechanism: increasing the relative joint-torque contribution changes the cost ordering of feasible sequences and favors coordinated motion with lower direct pitch-torque demand. Final error was not strictly monotonic. The smallest error, approximately 0.0374 m, occurred at $(\alpha_1, \alpha_2) = (1.4, 1.0)$, whereas the lowest Stage 1 pitch-torque RMS, approximately 0.0566 N·m, occurred at $(1.4, 1.4)$. Coupling guidance therefore primarily changed actuator allocation rather than directly optimizing geometric accuracy.

The nominal pair $(\alpha_1, \alpha_2) = (1,1)$ yielded a final error of approximately 0.0395 m and a Stage 1 pitch-torque RMS of approximately 0.0797 N·m. This pair lies in a balanced region rather than at an extreme operating point selected for a single metric. The maximum Stage 1 saturation ratio ρ_{sat} among all tested pairs was only about 1.41%. The reduction in $\tau^{(1)}_{\theta,\mathrm{RMS}}$ therefore arose mainly from redistribution within the feasible set rather than prolonged operation on an actuator boundary.

Fig. 7 relates final error to Stage 1 pitch-torque RMS $\tau^{(1)}_{\theta,\mathrm{RMS}}$.

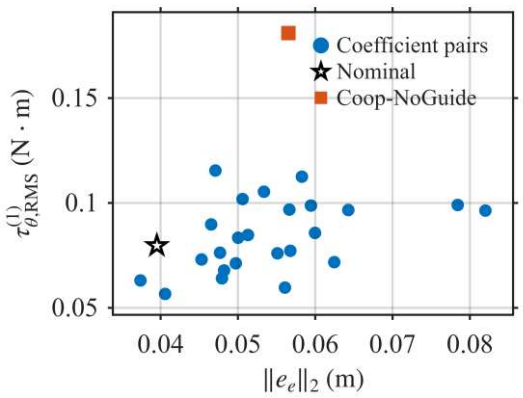
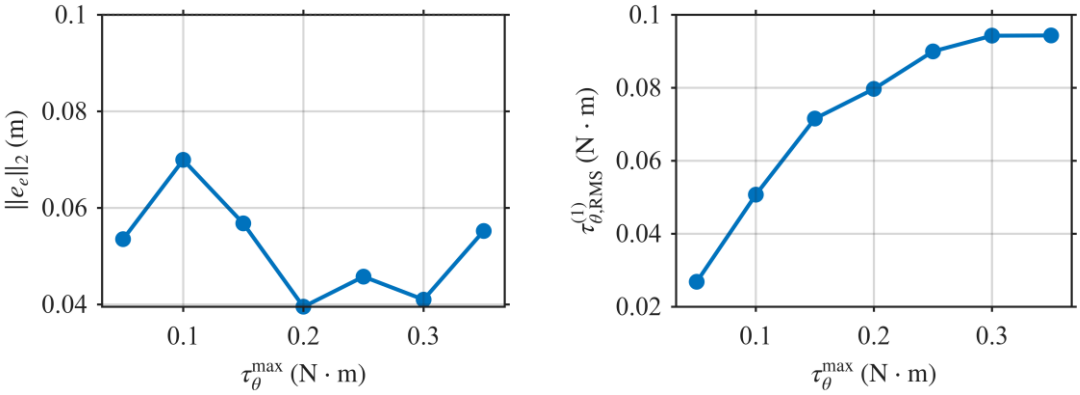


Fig. 7 Relationship between final end-effector error and Stage 1 pitch-actuator demand

The tested coefficient pairs yielded substantially lower pitch-torque demand than Coop-NoGuide, while the nominal pair also maintained a small final error. Pitch unloading therefore persisted across the tested coefficient neighborhood rather than occurring only at a single empirically selected point.

The admissible direct pitch-torque limit $\tau_\theta^{\max}$ was then varied to examine task completion and allocation as direct pitch authority decreased. All other parameters were fixed:

$$\tau_\theta^{\max} \in \{0.05, 0.10, 0.15, 0.20, 0.25, 0.3, 0.35\} \ \mathrm{N\cdot m} \tag{67}$$



(a) Final end-effector error **(b) Stage 1 pitch-torque RMS**

Fig. 8 Final end-effector error and pitch-actuator demand under different direct pitch-torque limits

Fig. 8 shows final error and Stage 1 pitch-torque RMS under each direct pitch-torque limit. Coop-Guided completed the prescribed task for all tested limits, with final errors of 0.0395–0.0699 m. At the smallest tested limit, $\tau_\theta^{\max} = 0.05$ N·m, the final error was approximately 0.0535 m. The switch to Stage 2 occurred within 12.78–12.84 s across the scan, without a pronounced delay as $\tau_\theta^{\max}$ decreased. Within the tested range, the control strategy adjusted base–joint motion and input allocation as pitch authority changed. It did not rely on a large direct pitch torque for the long-range approach.

This result is consistent with the nominal comparison and ablation. As the admissible direct pitch-torque range narrowed, the available base attitude authority decreased. Because the manipulator degrees of freedom remained in the unified prediction model, the optimizer could adjust the base–joint motion combination and complete the task throughout the tested range. The constraint scan therefore provides complementary evidence for the actuator-allocation interpretation.

Coop-Guided completed the prescribed task for every tested direct pitch-torque limit from 0.05 to 0.35 N·m. The value 0.05 N·m is only the smallest successful limit tested here. It is not a theoretical controllability threshold, a minimum feasible actuator capability, or a universal pitch-actuator design bound. Feasibility below this limit would depend on the initial condition, target, manipulator parameters, prediction horizon, and other state and input constraints.

Across the nominal comparison, ablation, and parameter scans, the observed direct pitch unloading was not confined to a single parameter combination. Coordination of all five degrees of freedom and coupling guidance jointly changed the base–joint motion and input sequences used to perform the task. Direct pitch-channel demand consequently remained low within the tested coefficient and actuator-authority ranges.

5. Conclusions

A coupling-guided two-stage NMPC method was developed for long-range target reaching by a quadrotor aerial manipulator under constrained direct pitch actuation. The control strategy uses a yaw-aligned five-degree-of-freedom longitudinal model to jointly optimize platform and manipulator motion. During the long-range approach, a coupling-guided allocation penalty and asymmetric actuator costs coordinate base motion, joint motion, and input allocation. Near the target, the control objective shifts to end-effector regulation with a base-drift penalty and a prediction-terminal contraction condition.

For the tested task, Coop-Guided maintained final positioning accuracy comparable to Fair Lock-3DOF while reducing the Stage 1 pitch-torque RMS by 51.1% and the accumulated absolute pitch torque by 62.2%. The ablation and common-window statistics indicate that the guidance term changes the optimizer's preference among feasible base–joint motion and input sequences rather than merely transferring pitch demand to the joints. The same unloading trend persisted across the tested guidance-coefficient combinations and direct pitch-torque limits.

These results indicate that base–manipulator dynamic coupling can provide an additional pathway for task realization under limited pitch authority within a unified predictive control strategy, rather than serving only as an internal disturbance to be suppressed. The findings are limited to the tested target-reaching task, the five-degree-of-freedom longitudinal model, and the investigated parameter and actuator-authority ranges. Future work will extend the framework to full three-dimensional motion, address model uncertainty, and validate the method on a physical aerial manipulation platform.

References

[1] Zhang, Y., Li, Z., Yu, Y., Fan, W., Xu, T., Li, D., Yu, X., Xu, B.: From aerial transportation to aerial interaction: A review. Chin. J. Aeronaut. Article 104442 (2026). https://doi.org/10.1016/j.cja.2026.104442

[2] Ollero, A., Tognon, M., Suarez, A., Lee, D., Franchi, A.: Past, present, and future of aerial robotic manipulators. IEEE Trans. Robot. 38(1), 626–645 (2022). https://doi.org/10.1109/TRO.2021.3084395

[3] Tullu, A., Jung, S.: Design, modeling, and control of Gimbal-robotic-arm for a fixed-wing unmanned aerial vehicle with vertical takeoff and landing capabilities. Chin. J. Aeronaut. Article 104381 (2026). https://doi.org/10.1016/j.cja.2026.104381

[4] Xu, W.-H., Cao, L.-J., Zhong, C.-L.: Review of aerial manipulator and its control. Int. J. Robot. Control Syst. 1(3), 308–325 (2021). https://doi.org/10.31763/ijrcs.v1i3.363

[5] Ruggiero, F., Lippiello, V., Ollero, A.: Aerial manipulation: A literature review. IEEE Robot. Autom. Lett. 3(3), 1957–1964 (2018). https://doi.org/10.1109/LRA.2018.2808541

[6] Samadikhoshkho, Z., Lipsett, M.: Decoupled control design of aerial manipulation systems for vegetation sampling application. Drones 7(2), 110 (2023). https://doi.org/10.3390/drones7020110

[7] Yadav, R.D., Dantu, S., Pan, W., Sun, S., Roy, S., Baldi, S.: Modular adaptive aerial manipulation under unknown dynamic coupling forces. IEEE/ASME Trans. Mechatron. 30(4), 2688–2698 (2025). https://doi.org/10.1109/TMECH.2024.3457806

[8] Mimmo, N., Macchelli, A., Naldi, R., Marconi, L.: Robust motion control of aerial manipulators. Annu. Rev. Control 49, 230–238 (2020). https://doi.org/10.1016/j.arcontrol.2020.04.006

[9] Li, H., Li, Z., Song, F., Yu, X., Yang, X., Rodríguez-Andina, J.J.: Finite-time fast adaptive backstepping attitude control for aerial manipulators based on variable coupling disturbance compensation. IEEE Trans. Ind. Electron. 71(11), 14730–14739 (2024). https://doi.org/10.1109/TIE.2024.3383036

[10] Li, Z., Li, H., Xu, Q., Yu, X., Basin, M.V.: Coupling disturbance modeling and compensation for aerial manipulator in highly dynamic motion. IEEE Trans. Cybern. 55(1), 124–135 (2025). https://doi.org/10.1109/TCYB.2024.3476124

[11] Samadikhoshkho, Z., Ghorbani, S., Janabi-Sharifi, F., Zareinia, K.: Nonlinear control of aerial manipulation systems. Aerosp. Sci. Technol. 104, 105945 (2020). https://doi.org/10.1016/j.ast.2020.105945

[12] Choi, J., Lee, D., Lee, I.: Model predictive control of aerial manipulation via hardware-in-the-loop simulation. IEEE Access 13, 197043–197052 (2025). https://doi.org/10.1109/ACCESS.2025.3634235

[13] Xu, B., Dai, Y., Shi, Y.: Stability-guaranteed model predictive control for UAVs: A survey of the Lyapunov-based framework. Chin. J. Aeronaut. 39(7), 3–21 (2026). https://doi.org/10.1016/j.cja.2026.104117

[14] Feng, R., Liu, Z., Luo, K., Liu, Y., Wang, X.: Uncertainty-aware model predictive control with learning-based residual dynamics compensation for agile quadrotor trajectory tracking. Chin. J. Aeronaut. Article 104424 (2026). https://doi.org/10.1016/j.cja.2026.104424

[15] Cao, H., Wu, Y., Wang, L.: Adaptive NN motion control and predictive coordinate planning for aerial manipulators. Aerosp. Sci. Technol. 126, 107607 (2022). https://doi.org/10.1016/j.ast.2022.107607

[16] Su, Y., Li, J., Jiao, Z., Wang, M., Chu, C., Li, H., Zhu, Y., Liu, H.: Sequential manipulation planning for over-actuated unmanned aerial manipulators. In: Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 6905–6911. IEEE (2023). https://doi.org/10.1109/IROS55552.2023.10341441

[17] Deng, W., Chen, H., Ye, B., Chen, H., Li, Z., Lyu, X.: Whole-body integrated motion planning for aerial manipulators. IEEE Trans. Robot. 41, 6661–6679 (2025). https://doi.org/10.1109/TRO.2025.3626619

[18] Pasetto, A., Vyas, Y., Cocuzza, S.: Zero reaction torque trajectory tracking of an aerial manipulator through Extended Generalized Jacobian. Appl. Sci. 12(23), 12254 (2022). https://doi.org/10.3390/app122312254

[19] Laribi, N., Mokhtari, M.R., Benallegue, A., El-Hadri, A., Benallegue, M.: QP-based control of an underactuated aerial manipulator under constraints. Aerosp. Sci. Technol. 177, 112845 (2026). https://doi.org/10.1016/j.ast.2026.112845

[20] Pose, C., Giribet, J., Mas, I.: Adaptive center-of-mass relocation for aerial manipulator fault tolerance. IEEE Robot. Autom. Lett. 7(2), 5583–5590 (2022). https://doi.org/10.1109/LRA.2022.3154040

[21] Li, H., Li, Z., Wu, T., Dong, C., Xu, Q., Yang, Y., Yu, X.: A novel active anti-disturbance control strategy for unmanned aerial

manipulator based on variable coupling disturbance compensation. Electronics 13(8), 1477 (2024). https://doi.org/10.3390/electronics13081477

Statements and Declarations

CRediT authorship contribution statement

Qingling Lin: Formal analysis, Software, Validation, Visualization, Writing – original draft. Gefei Shi: Conceptualization, Methodology, Funding acquisition, review & editing. Zhigang. Wu: Conceptualization, Methodology, Validation, Visualization, Funding acquisition, Writing – review & editing.

Declaration of competing interest

None.

Data availability

Data will be made available on request.

Acknowledgments

This work was partially funded by the Basic Science Center Program of the National Natural Science Foundation of China (62388101), State Key Laboratory of Mechanics and Control for Mechanical Structures (Nanjing University of Aeronautics and astronautics) (Grant No. MCAS-E-0226G01).