

Aerial Manipulation with Contact-Aware Onboard Perception and Hybrid Control

Yuanzhu Zhan^{*1}, Yufei Jiang^{*1}, Muqing Cao², Junyi Geng¹

Abstract—Aerial manipulation (AM) promises to move Unmanned Aerial Vehicles (UAVs) beyond passive inspection to contact-rich tasks such as grasping, assembly, and in-situ maintenance. Most prior AM demonstrations rely on external motion capture (MoCap) and emphasize position control for coarse interactions, limiting deployability. We present a fully onboard perception–control pipeline for contact-rich AM that achieves accurate motion tracking and regulated contact wrenches without MoCap. The main components are (1) an augmented visual–inertial odometry (VIO) estimator with contact-consistency factors that activate only during interaction, tightening uncertainty around the contact frame and reducing drift, and (2) image-based visual servoing (IBVS) to mitigate perception–control coupling, together with a hybrid force–motion controller that regulates contact wrenches and lateral motion for stable contact. Experiments show that our approach closes the perception-to-wrench loop using only onboard sensing, yielding an velocity estimation improvement of 66.01% at contact, reliable target approach, and stable force holding—pointing toward deployable, in-the-wild aerial manipulation.

I. INTRODUCTION

Despite the rapid growth, most Unmanned Aerial Vehicles (UAVs) today are limited to passive tasks like visual inspection, surveillance, and remote sensing [1], [2], where the UAVs avoid collision without physical interaction. Aerial manipulation (AM), where UAVs equipped with manipulators performing contact-based tasks such as grasping, assembly, transporting, and manipulating objects at height, will unlock transformative applications like infrastructure maintenance (e.g., bridge repair and relic restoration), precision agriculture, and contact inspection in manufacturing. Such capabilities can reduce human risk, enhance efficiency in labor-intensive sectors, and expand access to extreme or hazardous environments. The key challenge lies in the dynamical coupling between the UAVs and their interactions with other entities, where floating base body motion affects the end-effector and vice versa, as well as the inherent instability of multirotor platforms.

Most aerial manipulation demonstrations to date have been conducted in controlled lab settings—e.g., contact-based inspection [3], [4], door opening [5], valve turning [6], cleaning [7], and letter drawing [8]. These studies typically rely on external motion capture (MoCap) to provide near-ground-truth localization and state estimates. That assumption breaks

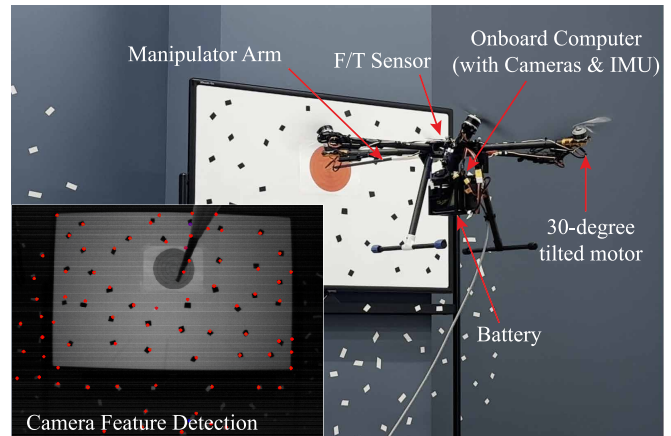


Fig. 1. Our fully-actuated aerial manipulator performs contact inspection using solely onboard sensing for motion tracking and force regulation. The inset shows the camera-based feature detection of our contact-aware visual inertial odometry.

down in real-world deployments, where external tracking is unavailable and GPS can be unreliable or denied especially near large infrastructure, indicating the necessity of fully onboard perception and control.

Emerging efforts have begun to investigate aerial manipulation in the wild with outdoor demonstrations, including GPS/RTK-assisted sensor retrieval [9], [10], total-station-based inspection [11], and vision-based coarse pick-and-place of large objects [12]. However, these systems typically address simple or low-precision tasks dominated by position control, or rely on unconventional mechanical designs, e.g., soft grippers [13] and compliant continuum manipulators [14] for position-based grasping. To our knowledge, they rarely model or regulate contact wrenches (force/torque), which is critical for contact-rich aerial manipulation tasks such as aerial contact inspection and peg-in-hole for maintenance, where maintaining specified wrenches is essential.

We aim to achieve precise aerial manipulation tasks in real world using only onboard sensing for maintaining accurate motion tracking and regulated contact wrenches. This is challenging for several reasons. First, the localization accuracy of standard vision-based Simultaneous Localization and Mapping (SLAM) / Visual Inertial Odometry (VIO) is often insufficient for simultaneous motion and force/torque control, especially near infrastructure where features are sparse or occluded. In addition, due to the tight dynamic coupling between the UAV body and the manipulator, perception and control are deeply intertwined: estimation errors can destabilize interaction control, while suboptimal control further degrades perception. These factors motivate a jointly

* Equal contribution.

¹ Department of Aerospace Engineering, Pennsylvania State University, University Park, PA, 16802, USA. {yvz6008, yufei-jiang, jgeng}@psu.edu

² Robotics Institute, Carnegie Mellon University, Pittsburgh PA 15213, USA. muqingc@andrew.cmu.edu

designed perception–control framework.

To address these challenges, this paper develops a contact-aware, fully onboard perception–control pipeline for aerial manipulation. Specifically, we augment the state-of-the-art VIO algorithm [15], [16], [17], [18] with contact consistency factors that activate only during interaction to stabilize estimates at contact. This tightens the uncertainty around the contact frame and reducing drift that would otherwise degrade interaction quality. To mitigate perception–control coupling issue of aerial manipulation, we adopt visual servoing to co-design perception and control. Image-based visual servoing (IBVS) directly uses image-space feedback to compute commands and thus does not require full UAV pose estimation [19]. We thus use only the VIO-provided body velocity for control feedback. Finally, a hybrid force–motion controller regulates contact wrenches along the surface normal while tracking lateral motion, maintaining both target pose and desired wrench during manipulation. In summary, the main contributions of this work are:

- We propose a contact-aware, *onboard* perception–control pipeline for contact-rich aerial manipulation that achieves both precise motion tracking and wrench regulation *without external motion capture*.
- We augment factor-graph VIO with *contact-consistency* constraints that activate only during interaction to improve estimation performance at contact.
- We develop an image-based visual servoing scheme that couples perception and control, together with a hybrid force–motion controller that regulates contact wrenches and lateral motion.
- We present simulation and real world experiments to demonstrate the effectiveness of accurate motion tracking and stable force holding under visual guidance.

II. RELATED WORKS

A. Aerial Manipulation

Research on aerial manipulation spans diverse tasks, such as aerial grasping [20], aerial docking [21], drilling [22], screwing [23], and pushing a target [24]. However, most demonstrations rely on external motion-capture (MoCap) for localization and are conducted indoors. Prior work has largely emphasized mechanical design [25], [9], [26] and dynamic analysis/control [27], [28], [29], with comparatively less attention to onboard perception [30]. Although emerging efforts push aerial manipulation into the wild with outdoor demonstrations, they often depend on external infrastructure such as GPS/RTK [9], [10] or total stations [11], and only a few use onboard sensing with standard VIO/SLAM for coarse pick-and-place of large objects [30]. These typically employ position-only control without explicit wrench regulation during physical interaction. In contrast, our goal is precise, MoCap-free aerial manipulation using only onboard sensing to control both motion and wrench.

B. Onboard Perception for Aerial Robots

Onboard perception for UAVs is well studied across platforms and tasks ranging from inertial+GPS navigation [31] to

LiDAR/sonar mapping [32] and, more recently, lightweight vision pipelines for tracking [33], landing [34], and inspection [35]. In aerial manipulation, however, onboard perception remains comparatively less mature. Existing aerial manipulation systems often employ VIO/SLAM (sometimes aided by fiducial markers) for approach behaviors such as grasping [13], perching [36], or package delivery [37]. However, contact is typically treated as a disturbance where texture loss, self-occlusion, and low-speed motion weaken visual constraints when accuracy is most critical. Consequently, MoCap-free demonstrations tend to emphasize coarse placement or light interaction and rely on position-only control without explicit wrench regulation [30], [13]. Outside the aerial domain, contact-aided estimation has improved observability by injecting no-slip/contact constraints into factor-graph formulations for legged or mobile manipulation [38], [39]. Inspired by this principle, we integrate contact-consistency into a VIO that activates only during interaction to stabilize pose/velocity at contact. To connect perception to action, visual servoing provides an interface that has been widely used for UAV tasks [40], [41]. Yet, prior aerial manipulation pipelines rarely close the full perception-to-wrench loop without MoCap; those that do are typically validated only in simulation [42]. In contrast, our work explicitly bridges targeting and contact using only onboard sensing leveraging velocity estimates from our contact-aware VIO to drive visual servoing and hybrid force–motion control.

III. SYSTEM OVERVIEW

A. System Design

As illustrated in Figure 1, our customized aerial platform is a hexarotor (Tarot T960) with six rotors, each tilted by 30 degrees in alternating left and right directions. Every arm carries a brushless motor driving a 12-inch propeller. Leveraging its fully actuated nature, we mount a zero degree-of-freedom (DOF) manipulator arm to the front of the vehicle without the need for additional complex mechanical components. A 6-axis force/torque (F/T) sensor ATI Gamma is installed at the base of the manipulator for interaction forces and moments measurement.

As for the avionics, we use the compact ModalAI VOXL 2 Flight Deck¹, which integrates an autopilot running PX4 firmware, an onboard computer, an inertial measurement unit (IMU), a barometer, and multiple optical sensors into a single unit. The platform provides a stereo pair for VIO and an RGB camera for visual servoing, meeting our sensing requirements.

B. Notations

We define five reference frames: the world inertial frame \mathcal{W} , the vehicle body frame \mathcal{B} , the camera frame \mathcal{O} , the F/T sensor frame \mathcal{S} , and the contact frame \mathcal{C} .

The inertial frame \mathcal{W} is a north-east-down frame whose origin is located at the UAV’s takeoff position. The body

¹<https://www.modalai.com/>

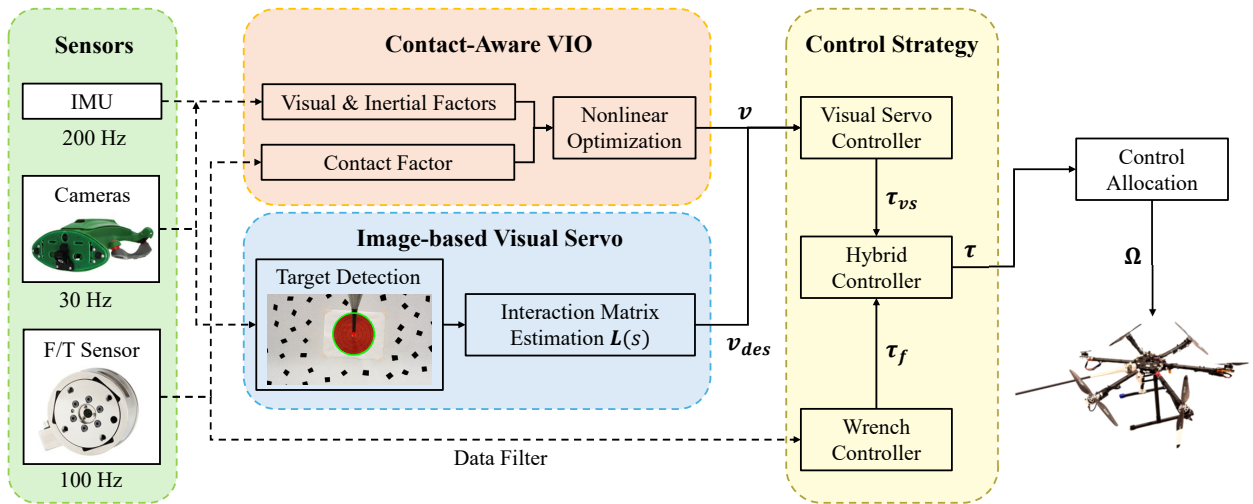


Fig. 2. Overview of the proposed pipeline. A contact-aware VIO module fuses visual, inertial, and force/torque measurements to provide state estimates and activates contact factors during interaction. An image-based visual servo (IBVS) drives the vehicle during the approach phase using VIO-estimated body velocity. Upon contact, a hybrid motion–wrench control strategy regulates contact forces and tracks lateral motion. Finally, a control allocation module maps the commanded wrenches into rotor speeds Ω for execution on the fully-actuated hexarotor platform.

frame is denoted as $\mathcal{B} = \{O_{\mathcal{B}}, \hat{x}_b, \hat{y}_b, \hat{z}_b\}$, where $O_{\mathcal{B}}$ represents the vehicle’s center of mass, and $\hat{x}_b, \hat{y}_b, \hat{z}_b$ are unit vectors pointing forward-left-up relative to the vehicle, respectively. The camera frame is defined as $\mathcal{O} = \{O_{\mathcal{O}}, \hat{x}_o, \hat{y}_o, \hat{z}_o\}$, where $O_{\mathcal{O}}$ is the optical center of the camera, and $\hat{x}_o, \hat{y}_o, \hat{z}_o$ are unit vectors pointing right-down-forward relative to the camera. The F/T sensor frame is defined as $\mathcal{S} = \{O_{\mathcal{S}}, \hat{x}_s, \hat{y}_s, \hat{z}_s\}$, whose $O_{\mathcal{S}}$ is at the base of the manipulator where the F/T sensor is mounted, and $\hat{x}_s, \hat{y}_s, \hat{z}_s$ are unit vectors aligned with the left, down, and up (left-handed coordinates) directions of the sensor. The contact frame is given by $\mathcal{C} = \{O_{\mathcal{C}}, \hat{x}_c, \hat{y}_c, \hat{z}_c\}$, where $O_{\mathcal{C}}$ corresponds to the end-effector tip at the contact point, and $\hat{x}_c, \hat{y}_c, \hat{z}_c$ are unit vectors aligned with the surface normal (pointing inward), left, and up directions of the contact surface.

Furthermore, $\mathbf{R}_a^b \in \text{SO}(3)$ and $\mathbf{T}_a^b \in \text{SE}(3)$ define the rotation and transformation matrix for vector from frame b to frame a . The right superscripts, e.g., $\mathbf{v}^{\mathcal{W}}$, indicates that the vector \mathbf{v} is expressed in frame \mathcal{W} . $\mathbf{0}_{a \times b}$ denotes a zero vector of dimension $a \times b$.

C. System Dynamics

Our platform is a fully-actuated aerial manipulator. Unlike conventional coplanar multirotor whose thrust is fixed along the body z -axis and must tilt the entire vehicle to generate lateral forces, a fully-actuated design can produce a full 6 DOF wrench and thus regulate translation and orientation independently. This decoupling is well suited to contact-rich tasks such as aerial contact inspection.

The dynamics of our hexarotor UAV can be modeled using Euler-Lagrangian method as following:

$$\mathbf{M}\dot{\mathbf{v}} + \mathbf{C}\mathbf{v} + \mathbf{G} = \boldsymbol{\tau} + \boldsymbol{\tau}_{ext} \quad (1)$$

with inertia matrix $\mathbf{M} \in \mathbb{R}^{6 \times 6}$, the centrifugal and Coriolis matrix $\mathbf{C} \in \mathbb{R}^{6 \times 6}$, the gravity vector $\mathbf{G} \in \mathbb{R}^6$. $\mathbf{v} =$

$[\mathbf{v}^{\top}, \boldsymbol{\omega}^{\top}]^{\top}$ and $\dot{\mathbf{v}}$ represent the system twist and acceleration respectively. $\boldsymbol{\tau} \in \mathbb{R}^6$ is the control wrench from the actuators, while $\boldsymbol{\tau}_{ext} \in \mathbb{R}^6$ is other external wrench (e.g. wind or contact).

D. Method Overview

We target aerial manipulation tasks that require precise motion tracking and regulated contact wrenches. The representative scenario is approaching a surface and sustaining a prescribed normal force, such as in aerial contact inspection or peg-in-hole maintenance. We focus on the problem to the approach and interaction phases because long-range, free-flight guidance can be handled by other standard and mature methods and is outside the scope of this work. The proposed pipeline as shown in Figure. 2 has three key components: (1) A contact-aware VIO processes stereo imagery to produce vehicle/end-effector state estimates, activating contact factors only during interaction. (2) During approach, image-based visual servoing (IBVS) drives the platform using only the VIO-estimated body velocity. (3) Upon contact, a hybrid force–motion controller regulates the normal wrench while IBVS continues to track lateral motion, closing the perception-to-wrench loop with onboard sensing.

IV. CONTACT-AWARE VISUAL-INERTIAL ODOMETRY

Our method builds upon VINS-Fusion [15], [16], where visual–inertial state estimation is expressed within a factor graph framework, and the system state is recovered through joint optimization of camera reprojection errors and IMU pre-integration constraints. While this formulation achieves high accuracy in general scenarios, it may suffer from drift when external interactions occur. To mitigate this issue, we extend it by introducing contact factors that leverage robot–environment interactions. In particular, when the system is in contact with a wall, the relative motion at the contact point is constrained, and this information is encoded




Fig. 3. Pose graph of the proposed contact-aware VIO

as an additional factor in the graph as shown in Figure 3. By incorporating such constraints, the estimator gains improved robustness and reduced drift, especially in environments with limited visual features or long-term operation.

A. Visual-inertial Odometry

We formulate visual-inertial odometry with contact constraints as a nonlinear optimization problem over a sliding window of states. The full state vector is defined as

$$\mathcal{X} = [\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{x}_c^B, \gamma_0, \gamma_1, \dots, \gamma_m], \quad (2)$$

where each aerial manipulator state \mathbf{x}_k at time $k \in [0, n]$ includes the position, velocity, orientation, and bias parameters of the accelerometer and gyroscope, and we assume that the body frame of the aerial manipulator is aligned with the IMU frame.

$$\mathbf{x}_k = [\mathbf{p}_k^W, \mathbf{v}_k^W, \mathbf{q}_k^W, \mathbf{b}_a, \mathbf{b}_g] \quad (3)$$

and $\mathbf{x}_c^B = [\mathbf{p}_c^B, \mathbf{q}_c^B]$ denotes the translation and orientation of camera in the body frame. The variables γ_i represent the inverse depths of visual landmarks.

Given this state representation, the state estimation problem is expressed as minimizing the sum of residuals from different sensing modalities:

$$\min_{\mathcal{X}} \|r_p - H_p \mathcal{X}\|^2 + \sum_{(i,j) \in \mathcal{C}} \|r_{\text{visual}}^{ij}(\mathcal{X})\|^2 + \sum_{k \in \mathcal{I}} \|r_{\text{IMU}}^k(\mathcal{X})\|^2. \quad (4)$$

r_{visual}^{ij} denotes the reprojection residual of feature j observed at time i , r_{IMU}^k denotes the IMU pre-integration residual between consecutive states. This joint optimization ensures consistency between visual and inertial measurements and provides accurate state estimation in general scenarios. The visual residuals and IMU residuals follow the standard formulation in [15], and we refer readers to that paper for more detailed derivations.

B. Contact Factor

To further enhance robustness in scenarios involving robot-environment interactions, we introduce a contact factor into the factor graph. When the system establishes contact with a wall, the relative motion at the contact point can be assumed to follow physical constraints, which are encoded as additional residuals in the optimization. Formally, the

extended objective becomes

$$\min_{\mathcal{X}} \|r_p - H_p \mathcal{X}\|^2 + \sum_{(i,j) \in \mathcal{V}} \|r_{\text{visual}}^{ij}(\mathcal{X})\|^2 + \sum_{k \in \mathcal{I}} \|r_{\text{IMU}}^k(\mathcal{X})\|^2 + \sum_{k \in \mathcal{F}} \|r_{\text{contact}}^k(\mathcal{X})\|^2, \quad (5)$$

where r_{contact}^k enforces the constraint induced by wall contact at time k . Here, \mathcal{V} denotes the set of all visual feature observations, \mathcal{I} denotes the set of all IMU measurement intervals, and \mathcal{F} denotes the set of all contact events.

The contact factor reduces drift and improves estimation accuracy by anchoring the states to the environment during contact events, especially under feature-sparse or long-term operation.

Concretely, the contact event is triggered by a force threshold. After we switch to force holding mode, the robot is constrained along the wall normal direction. Between two consecutive states, the signed distance to the wall is preserved, and the velocity projected onto the wall normal is zero. These constraints are expressed as

$$r_{\text{contact}}^k(\mathbf{x}) = \begin{bmatrix} \mathbf{n}^\top (\mathbf{p}_{k+1}^W - \mathbf{p}_k^W) \\ \mathbf{n}^\top \mathbf{v}_k^W \end{bmatrix}, \quad (6)$$

where \mathbf{n} denotes the wall normal, \mathbf{p}_k and \mathbf{p}_{k+1} are positions at consecutive time steps, and \mathbf{v}_k is the velocity at time k in the world frame. By jointly constraining both position and velocity along the contact normal, the factor further stabilizes the state estimation in challenging environments.

To further improve robustness, the covariance of the contact factor is adaptively adjusted according to the force sensor readings. Intuitively, when the contact force is stable and large, the system is confident that the robot is in firm contact with the environment, and the corresponding constraint should be enforced more strongly. Conversely, if the force sensor signal is noisy or unstable, this indicates that the contact is not reliable, and the constraint should be relaxed. Formally, the information matrix of the contact residual is defined as

$$\mathbf{P}_{\text{contact}} = \alpha \frac{1}{N} \sum_{i=1}^N \left(F_i - \frac{1}{N} \sum_{j=1}^N F_j \right)^2, \quad (7)$$

where F_i denotes the measured contact force along the constrained direction at the i -th sample within a sliding window of length N , and α is a scaling coefficient. This variance-based formulation captures the stability of the contact interaction: stable and consistent contact forces yield a small variance and thus a stronger constraint, while noisy or fluctuating forces produce a larger variance, which automatically weakens the influence of the contact factor and prevents the introduction of spurious constraints.

In the optimization, the adaptive weighting enters the objective through the Mahalanobis norm

$$\|r_{\text{contact}}^k(\mathcal{X})\|^2 = (r_{\text{contact}}^k(\mathcal{X}))^\top \mathbf{P}_{\text{contact}}^{-1} r_{\text{contact}}^k(\mathcal{X}). \quad (8)$$

V. HYBRID MOTION AND FORCE CONTROL WITH IMAGE-BASED VISUAL SERVO

Thanks to the fully-actuated property, we design the control input τ^B via feedback linearization as the combination of a motion controller τ_{vs}^B (driven primarily by visual servoing) and a wrench controller τ_f^C based on [8], [43], [42]:

$$\tau^B = \mathbf{R}_B^C((\mathbf{I}_{6 \times 6} - \mathbf{\Lambda})\mathbf{R}_C^B\tau_{vs}^B + \mathbf{\Lambda}\tau_f^C) \quad (9)$$

where $\mathbf{\Lambda} = \text{blockdiag}[\lambda(d), 0, 0, 0, 0, 0] \in \mathbb{R}^{6 \times 6}$, $\lambda(d) \in \{0, 1\}$ is the selection matrix that specifies which DoFs are executed as direct wrench commands, while the complementary subspace is handled by motion control. $\lambda(d)$ is determined by the camera-derived depth d to the contact surface, smoothly transitioning from motion to force regulation as contact is approached. $\lambda(d) = 0$ corresponds to pure motion control driven by visual servoing, while $\lambda(d) = 1$ activates force control along the surface normal direction. To enhance a smooth transition, we design $\lambda(d)$ as a confidence factor that increases progressively with the measured depth d :

$$\lambda(d) = \begin{cases} 1, & d \leq d_{min}, \\ \frac{1}{2} \left(1 + \cos \left(\frac{d - d_{min}}{d_{max} - d_{min}} \pi \right) \right), & d_{min} \leq d \leq d_{max}, \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

A. Visual Servoing Strategy

We use the standard image based visual servo [19] with the velocity estimate from Sec. IV to guide the motion of the end-effector to align with the target on the surface.

Let us consider the real and desired image feature as \mathbf{s} , \mathbf{s}^* , respectively. The aim of IBVS is to minimize feature error $e_s = \mathbf{s} - \mathbf{s}^*$. For typical aerial contact inspection with motionless target, \mathbf{s}^* is constant, and changes in \mathbf{s} depend only on camera motion. Then, the first-order feature error dynamics can be represented as:

$$\dot{e}_s = \dot{\mathbf{s}} - \dot{\mathbf{s}}^* = \mathbf{L}(\mathbf{s})\mathbf{v}_c^O - \dot{\mathbf{s}}^* = \mathbf{L}(\mathbf{s})\mathbf{v}_c^O \quad (11)$$

where $\mathbf{v}_c \in \mathbb{R}^6$ is the spatial velocity of the camera (camera twist), including the instantaneous linear and angular components. $\mathbf{L}(\mathbf{s}) \in \mathbb{R}^{k \times 6}$ (k is the number of features) is the interaction matrix. We consider velocity input to the aerial manipulator controller. To impose an exponential decrease of the feature error, the camera twist can be designed as

$$\mathbf{v}_c^O = -\zeta \hat{\mathbf{L}}(\mathbf{s})^\dagger e_s \quad (12)$$

where $\hat{\mathbf{L}}(\mathbf{s})^\dagger$ denotes the estimated Moore-Penrose pseudo-inverse of $\mathbf{L}(\mathbf{s})$. $\zeta > 0$ is the convergence rate. Then, the velocity input to the UAV is $\mathbf{v}_{des}^B = \mathbf{R}_B^O\mathbf{v}_c^O$, and a standard PID controller can be designed to obtain τ_{vs}^B .

The key components are robust image-feature detection and accurate interaction matrix estimation. For the tasks such as peg-in-the hole, the target can be modeled as a circle. Standard circle detection and segmentation algorithm can be used. The detected image center $\mathbf{o}_c = [u, v]^\top$ and the circle's image area A encode lateral alignment and range (scale), respectively: the alignment error measures the lateral

offset between the end-effector tip and the target center, while the scaling error reflects distance to the target plane. We therefore choose the IBVS feature vector $\mathbf{s} = [\mathbf{o}_c^\top, r]^\top$, where r is the radius in pixel. Estimated depth is given by $d = \sqrt{f_x f_y \frac{A_r}{A}}$ with A the target area in pixels, A_r the known physical area of the circular target and f_x, f_y the camera focal lengths. $\mathbf{L}(\mathbf{s})$ can thus be obtained as:

$$\mathbf{L}(\mathbf{s}) = \begin{bmatrix} -\frac{f_x}{d} & 0 & \frac{u}{d} & \frac{uv}{f_x} & -\frac{f_x^2 + u^2}{f_x} & v \\ 0 & -\frac{f_y}{d} & \frac{v}{d} & \frac{f_x^2 + v^2}{f_y} & -\frac{uv}{f_y} & -u \\ 0 & 0 & \frac{r}{d} & -\frac{rv}{f_y} & \frac{ru}{f_x} & 0 \end{bmatrix} \quad (13)$$

B. Wrench Controller

We design an impedance controller to ensure reference force F_r tracking of the end-effector during the physical interaction with the target. The onboard F/T sensor measures the interaction force F . Then, the controller is designed as:

$$\begin{aligned} e_f &= F - F_r, \quad e_x = e_{s[3]} \\ F_f &= F_r - mm_d^{-1}(K_e e_x + D_e \dot{e}_x) - K_{f,p} e_f - K_{f,i} \int e_f dt \\ &= F_r - K_{e,p} e_x - K_{e,d} \dot{e}_x - K_{f,p} e_f - K_{f,i} \int e_f dt \end{aligned} \quad (14)$$

with m, m_d the actual and desired vehicle mass, respectively, $K_{e,p}, K_{e,d}$ are the normalized stiffness, and damping, respectively. $K_{f,p}, K_{f,i} > 0$ are tunable gains for force control. $e_{s[3]}$ is the third component of e_s , which reflecting the scaling error. Then, the control wrench $\tau_f^C = [F_f, \mathbf{0}_{1 \times 5}]^\top$ and τ_{vs}^B will be passed through (9) to get the total control wrench τ^B , which is finally distributed to each rotor through the control allocation process. Furthermore, leveraging the fully-actuated design, we command zero roll and pitch to keep the vehicle level during the manipulation task.

VI. EXPERIMENTS

We evaluate the proposed method in both simulation and real-world. In Gazebo simulation, we assess the robustness under varying levels of estimation error. Real-world experiments are then conducted to demonstrate contact-rich aerial manipulation using only onboard sensing.

A. Experiment Setup

In simulation, we model a fully actuated aerial manipulator in the Gazebo simulator following the UAV design in Sec. III-A. An RGB-D camera, configured with the same intrinsics/extrinsics as the VOXL2 Hi-Res optical sensor, is front-mounted at the same location on the vehicle. A 6-axis F/T sensor plugin is attached near the end-effector.

In real-world experiment, The aerial manipulator relies entirely on onboard sensing and computation as described in Sec. III-D. An OptiTrack motion-capture system provides ground-truth pose only for offline evaluation of the contact-aware VIO and is never used for feedback.

We consider a peg-in-hole task with force regulation as a representative aerial manipulation scenario of aerial contact




Fig. 4. Snapshots of real-world experiments. (a) Visual servoing initialized. (b) Approaching the target. (c) Force holding.

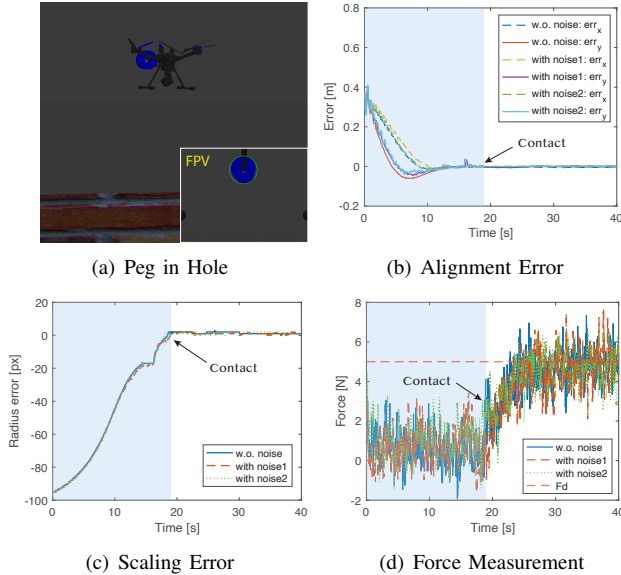


Fig. 5. Peg-in-hole experiment in simulation. (a) a snapshot of the contact moment. The first person view (FPV) is in the lower right corner. (b) Alignment error converted from the pixel plane. (c) Scaling error in pixel plane. (d) Filtered force measurement.

inspection and maintenance. The experimental procedure is as follows: First, the vehicle takes off and stabilizes in hover with the target in the camera’s field of view (FOV). Then the hybrid controller is activated and IBVS engaged. The color-based detector segments the target, and the resulting mask is used to extract visual features. Guided by our image-based visual servoing, the aerial manipulator approaches the surface, performs insertion, and regulates a specified normal force $F_d = 5$ N while maintaining lateral alignment.

B. Simulation Results

The experiment is conducted with a circular hole mounted on a vertical wall 1.5 m above the ground, with an inner diameter of 50 mm and an outer diameter of 140 mm; the aerial manipulator end-effector has a diameter of 12 mm. To verify the robustness of our proposed method, we repeat this experiment by gradually increasing levels of velocity estimation noise. The baseline uses ground-truth velocity. In the second condition, we inject zero-mean, uniformly distributed noise with bounds $(\pm 0.05, \pm 0.05, \pm 0.07)$ m/s (noise 1). In the third, we increase the bounds to $(\pm 0.10, \pm 0.10, \pm 0.14)$ m/s (noise 2). Figure 5 summarizes

force regulation and visual target tracking. In particular, Figures 5 (b)–(c) show IBVS feature tracking, where the lateral alignment error has been converted from pixels to meters using camera intrinsics and the target geometry, and the radius error is in pixel unit. We can see despite the injected velocity errors, the image feature estimation still converge. The aerial manipulator successfully performs the peg-in-hole task and holds reference $F_d = 5$ N force.

C. Real-world Experiments

We then evaluate the task in real-world experiments. To assess robustness, we set up a red circular wall target with a radius of 10 cm and four equally spaced concentric rings. Figure. 4 shows a sequence of flight snapshots (with an First person view (FPV) inset at the lower left). From $t = 0$ to 28.3s, IBVS guides the vehicle toward the vicinity of target and the lateral alignment error converges to ≈ 0 m. At $t = 28.3$ s, the end-effector made contact with the target, indicated by a sharp rise in the measured force. From 28.3 to 55.1 s, the hybrid force–motion controller maintains the commanded normal force while simultaneously regulating lateral motion, as shown in Figure. 6 (a)–(b). Thanks to the fully-actuated mechanism, the platform also maintains a near-zero body attitude (roll/pitch) during contact (Figure. 6(c)).

TABLE I
VELOCITY ESTIMATION ERROR DURING CONTACT (M/S) (\downarrow)

Method	RMSE	mean	max	std
VINS-Fusion	0.0356	0.0207	0.1174	0.0356
OpenVINS	0.0619	0.0493	0.2123	0.0619
Ours	0.0121	0.0089	0.0516	0.0117

We also evaluate the state estimation performance of our contact-aware VIO by comparing it against VINS-Fusion [15] and OpenVINS [44]. Figure 7 shows the estimated velocities along the x , y , and z axes, together with the velocity error in the contact direction (x -axis). While all three estimators generally follow the ground truth motion-capture (MoCap) measurements, significant differences appear in the contact direction. In particular, VINS-Fusion exhibits noticeable fluctuations during contact, resulting in large estimation errors, whereas our method remains stable and closely aligned with the MoCap ground truth. The error plot further highlights that the proposed approach suppresses high-frequency noise and reduces drift along the contact axis, leading to smoother and more accurate estimates.

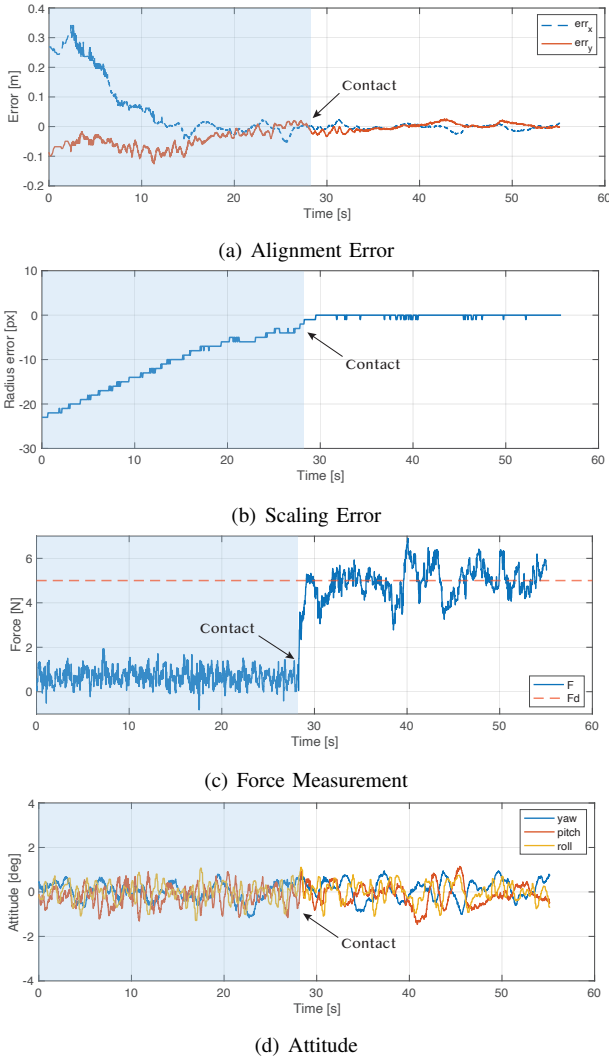


Fig. 6. Task evaluation in real world. (a) Alignment error converted from the pixel plane. (b) Scaling error in pixel plane (c) Filtered force measurement. (c) Flight attitude

A quantitative comparison is provided in Table. I, which reports the Root Mean Square Error (RMSE), mean, maximum, and standard deviation of the velocity error magnitude during contact. Our method achieves the lowest error across all metrics, with an RMSE of **0.0121 m/s** and mean error of only **0.0089 m/s**, significantly outperforming VINS-Fusion (RMSE 0.0356 m/s) and OpenVINS (RMSE 0.0619 m/s). These results clearly demonstrate that leveraging contact information substantially improves the robustness and accuracy of state estimation, particularly in challenging conditions where pure visual-inertial methods struggle.

VII. CONCLUSION

This paper developed a fully onboard perception-control pipeline for contact-rich aerial manipulation without external motion capture system. By adding contact-consistency factors to VIO and integrating IBVS with a hybrid force-motion controller, the system improves estimation during physical interaction and regulates both contact wrench and lateral motion. Experiments demonstrate accurate target approach,

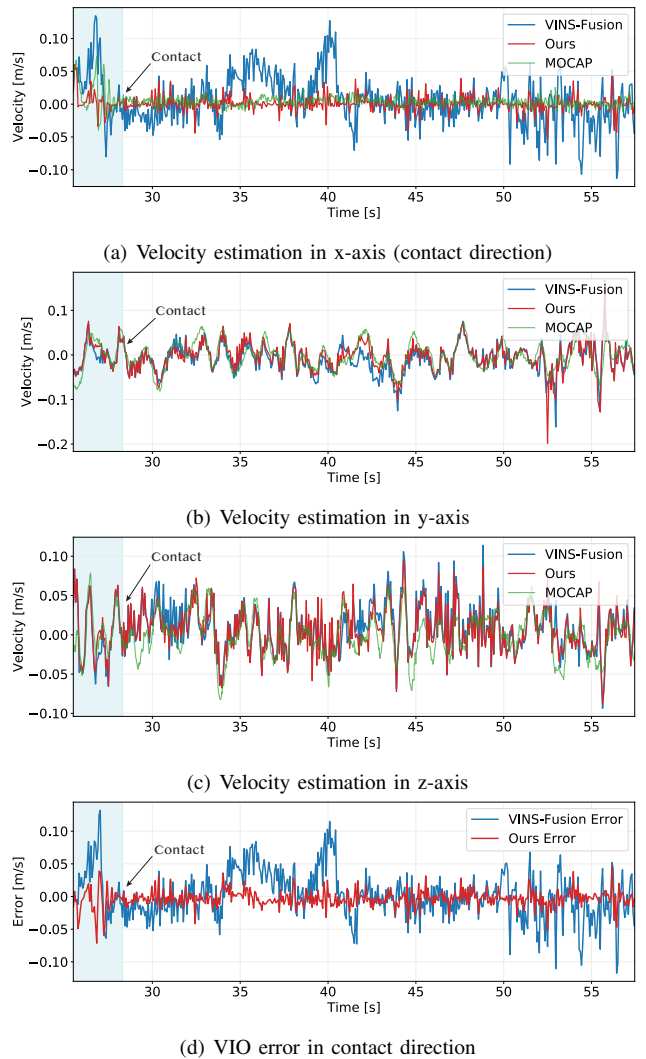


Fig. 7. Velocity estimation performance of VINS-Fusion and the proposed method, compared against ground-truth MoCap. (a)–(c) show velocity estimates along the x , y , and z axes, respectively, with the x -axis corresponding to the contact direction. (d) presents the velocity estimation error along the contact axis

stable force holding, and a 66.01% reduction in velocity estimation drift at contact with only onboard sensing, highlighting the potential for complex outdoor tasks. Future work will address online surface normal estimation and robustness to wind disturbance and feature-sparse environments.

REFERENCES

- [1] Y. Hu, J. Geng, C. Wang, J. Keller, and S. Scherer, “Off-policy evaluation with online adaptation for robot exploration in challenging environments,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 6, pp. 3780–3787, 2023.
- [2] Y. Qiu, C. Xu, Y. Chen, S. Zhao, J. Geng, and S. Scherer, “Airio: Learning inertial odometry with enhanced imu feature observability,” *IEEE Robotics and Automation Letters*, vol. 10, no. 9, pp. 9368–9375, 2025.
- [3] K. Bodie, M. Brunner, M. Pantic, S. Walser, P. Pfändler, U. Angst, R. Siegwart, and J. Nieto, “An omnidirectional aerial manipulation platform for contact-based inspection. arxiv 2019,” *arXiv preprint arXiv:1905.03502*.
- [4] X. Guo, G. He, M. Mousaei, J. Geng, G. Shi, and S. Scherer, “Aerial interaction with tactile sensing,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 1576–1582.

- [5] D. Lee, H. Seo, D. Kim, and H. J. Kim, "Aerial manipulation using model predictive control for opening a hinged door," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1237–1242.
- [6] M. Brunner, G. Rizzi, M. Studiger, R. Siegwart, and M. Togonon, "A planning-and-control framework for aerial manipulation of articulated objects," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 689–10 696, 2022.
- [7] M. Alkaddour, M. A. Jaradat, S. Tellab, N. A. Sherif, M. H. Alvi, L. Romdhane, and K. S. Hatamleh, "Novel design of lightweight aerial manipulator for solar panel cleaning applications," *IEEE Access*, vol. 11, pp. 111 178–111 199, 2023.
- [8] X. Guo, G. He, J. Xu, M. Mousaei, J. Geng, S. Scherer, and G. Shi, "Flying calligrapher: Contact-aware motion and force planning and control for aerial manipulation," *IEEE Robotics and Automation Letters (RA-L)*, vol. 9, no. 12, pp. 11 194–11 201, 2024.
- [9] S. Hamaza, I. Georgilas, G. Heredia, A. Ollero, and T. Richardson, "Design, modeling, and control of an aerial manipulator for placement and retrieval of sensors in the environment," *Journal of Field Robotics*, vol. 37, no. 7, pp. 1224–1245, 2020.
- [10] Z. Wen, Y. Liao, Y. Liu, X. Chen, and W. Meng, "Full body motion control for aerial manipulation based on critical dynamic constraints," *IEEE Robotics and Automation Letters*, 2025.
- [11] A. E. Jimenez-Cano, P. J. Sanchez-Cuevas, P. Grau, A. Ollero, and G. Heredia, "Contact-based bridge inspection multirotors: Design, modeling, and control considering the ceiling effect," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3561–3568, 2019.
- [12] E. Bauer, M. Blöchlinger, P. Strauch, A. Raayatsanati, C. Cavelti, and R. K. Katzschmann, "An open-source soft robotic platform for autonomous aerial manipulation in the wild," *arXiv preprint arXiv:2409.07662*, 2024.
- [13] S. Ubellacker, A. Ray, J. M. Bern, J. Strader, and L. Carlone, "High-speed aerial grasping using a soft drone with onboard perception," *npj Robotics*, vol. 2, no. 1, p. 5, 2024.
- [14] R. Peng, Y. Wang, M. Lu, and P. Lu, "A dexterous and compliant aerial continuum manipulator for cluttered and constrained environments," *Nature Communications*, vol. 16, no. 1, p. 889, 2025.
- [15] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [16] T. Qin and S. Shen, "Online temporal calibration for monocular visual-inertial systems," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3662–3669.
- [17] T. Qin, J. Pan, S. Cao, and S. Shen, "A general optimization-based framework for local odometry estimation with multiple sensors," 2019.
- [18] T. Qin, S. Cao, J. Pan, and S. Shen, "A general optimization-based framework for global pose estimation with multiple sensors," 2019.
- [19] F. Chaumette and S. Hutchinson, "Visual servo control. i. basic approaches," *IEEE Robotics & Automation Magazine*, vol. 13, no. 4, pp. 82–90, 2006.
- [20] D. Mellinger, Q. Lindsey, M. Shomin, and V. Kumar, "Design, modeling, estimation and control for aerial grasping and manipulation," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 2668–2673.
- [21] A. J. Choi, J. Park, and J.-H. Han, "Automated aerial docking system using onboard vision-based deep learning," *Journal of Aerospace Information Systems*, vol. 19, no. 6, pp. 421–436, 2022.
- [22] C. Ding, L. Lu, C. Wang, and C. Ding, "Design, sensing, and control of a novel uav platform for aerial drilling and screwing," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3176–3183, 2021.
- [23] M. Schuster, D. Bernstein, P. Reck, S. Hamaza, and M. Beiteltschmidt, "Automated aerial screwing with a fully actuated aerial manipulator," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3340–3347.
- [24] D. Lee, H. Seo, I. Jang, S. J. Lee, and H. J. Kim, "Aerial manipulator pushing a movable structure using a dob-based robust controller," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 723–730, 2020.
- [25] C. Lanegger, M. Ruggia, M. Togonon, L. Ott, and R. Siegwart, "Aerial layouting: Design and control of a compliant and actuated end-effector for precise in-flight marking on ceilings," *Proceedings of Robotics: Science and System XVIII*, p. p073, 2022.
- [26] M. Zhao, T. Anzai, F. Shi, X. Chen, K. Okada, and M. Inaba, "Design, modeling, and control of an aerial robot dragon: A dual-rotor-embedded multilink robot with the ability of multi-degree-of-freedom aerial transformation," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1176–1183, 2018.
- [27] M. Wang, Z. Chen, K. Guo, X. Yu, Y. Zhang, L. Guo, and W. Wang, "Millimeter-level pick and peg-in-hole task achieved by aerial manipulator," *IEEE Transactions on Robotics*, vol. 40, pp. 1242–1260, 2023.
- [28] W. Dong, Z. Ma, X. Sheng, and X. Zhu, "Centimeter-level aerial assembly achieved with manipulating condition inference and compliance," *IEEE/ASME Transactions on Mechatronics*, vol. 27, no. 3, pp. 1660–1671, 2021.
- [29] E. Cuniato, N. Lawrance, M. Togonon, and R. Siegwart, "Power-based safety layer for aerial vehicles in physical interaction using lyapunov exponents," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6774–6781, 2022.
- [30] M. Blöchlinger, Y. Toshimitsu, and R. K. Katzschmann, "Reliable aerial manipulation: Combining visual tracking with range sensing for robust grasping," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 7519–7525.
- [31] G. Hoffmann, H. Huang, S. Waslander, and C. Tomlin, "Quadrotor helicopter flight dynamics and control: Theory and experiment," in *AIAA guidance, navigation and control conference and exhibit*, 2007, p. 6461.
- [32] B. Zhou, Y. Zhang, X. Chen, and S. Shen, "Fuel: Fast uav exploration using incremental frontier structure and hierarchical planning," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 779–786, 2021.
- [33] S. Ghosh, J. Patrikar, B. Moon, M. M. Hamidi, and S. Scherer, "Airtrack: Onboard deep learning framework for long-range aircraft detection and tracking," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1277–1283.
- [34] S. Lin, M. A. Garratt, and A. J. Lambert, "Monocular vision-based real-time target recognition and tracking for autonomously landing an uav in a cluttered shipboard environment," *Autonomous Robots*, vol. 41, no. 4, pp. 881–901, 2017.
- [35] J. Xing, G. Cioffi, J. Hidalgo-Carrió, and D. Scaramuzza, "Autonomous power line inspection with drones via perception-aware mpc," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 1086–1093.
- [36] J. Mao, G. Li, S. Nogar, C. Kroninger, and G. Loianno, "Aggressive visual perching with quadrotors on inclined surfaces," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5242–5248.
- [37] G. Li, R. Ge, and G. Loianno, "Cooperative transportation of cable suspended payloads with mavs using monocular vision and inertial sensing," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5316–5323, 2021.
- [38] R. Hartley, J. Mangelson, L. Gan, M. G. Jadidi, J. M. Walls, R. M. Eustice, and J. W. Grizzle, "Legged robot state-estimation through combined forward kinematic and preintegrated contact factors," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4422–4429.
- [39] R. Buchanan, A. Röfer, J. Moura, A. Valada, and S. Vijayakumar, "Online estimation of articulated objects with factor graphs using vision and proprioceptive sensing," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16 111–16 117.
- [40] M. Xu, A. Hu, and H. Wang, "Image-based visual impedance force control for contact aerial manipulation," *IEEE Transactions on Automation Science and Engineering*, vol. 20, no. 1, pp. 518–527, 2022.
- [41] J. Byun, J. Kim, D. Eom, D. Lee, C. Kim, and H. J. Kim, "Image-based time-varying contact force control of aerial manipulator using robust impedance filter," *IEEE Robotics and Automation Letters*, vol. 9, no. 5, pp. 4854–4861, 2024.
- [42] G. He, Y. Jangir, J. Geng, M. Mousaei, D. Bai, and S. Scherer, "Image-based visual servo control for aerial manipulation using a fully-actuated uav," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 5042–5049.
- [43] K. Bodie, M. Brunner, M. Pantic, S. Walser, P. Pfändler, U. Angst, R. Siegwart, and J. Nieto, "Active interaction force control for contact-based inspection with a fully actuated aerial vehicle," *IEEE Transactions on Robotics*, vol. 37, no. 3, pp. 709–722, 2020.
- [44] P. Geneva, K. Eckenhoff, W. Lee, Y. Yang, and G. Huang, "Openvins: A research platform for visual-inertial estimation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4666–4672.