Atılım UAV Team Technical Design Report 2026

Atılım UAV Team
Atılım University
Ankara, Türkiye

***Abstract*—This report presents the design, development, and testing strategy of our autonomous unmanned aerial vehicle (UAV) developed for the 2026 Student Unmanned Aerial Systems (SUAS) competition, which targets disaster relief and emergency response missions. The aircraft carries the avionics and vision systems required for high-resolution aerial mapping and real-time target detection. The heavy-payload platform is built around a reinforced composite structure that provides structural integrity while keeping the airframe serviceable in the field. When a disaster victim is located, the system computes the target coordinates onboard and transmits them to the Ground Control Station, allowing rescue teams to be mobilized quickly. The aircraft then delivers supplies to the target area using a controlled-descent release system that requires no operator intervention. The platform has been validated through simulation and field testing covering flight behavior, navigation accuracy, onboard detection performance, and the integration of the mapping and payload subsystems.**

ACKNOWLEDGEMENTS

We would like to express our sincere gratitude to our academic advisor, Dr. Halis Kandaş, for his invaluable technical and logistical advice and continuous support. We also extend our thanks to all our professors, former and current team members, and mentors who pioneered the transition to autonomous systems and paved the way for our current work.

We owe a great debt of gratitude to our sponsors. Special thanks to ASELSAN, a leading defense industry company in Türkiye, whose financial backing enabled the optimal construction of our unmanned aerial vehicle (UAV). Furthermore, we deeply appreciate the Ministry of Youth and Sports for providing us with a dedicated facility to conduct our manufacturing, host team meetings, and organize social activities. We are also grateful to WupSoft and FSM Demirbaş for their assistance in procuring the fundamental components of our UAV. Finally, we would like to acknowledge Pegem Akademi, Çatkaya, ÖzcanSu, BallıKuyumcu Offroad, and Gezerler Construction for their generous financial and material support.

I. COMPETITION STRATEGY

A. Strategic Vision

Our strategic vision is to develop a reliable autonomous aerial platform capable of supporting disaster relief and emergency response operations while accomplishing all Student Unmanned Aerial Systems (SUAS) mission objectives with maximum precision and operational efficiency. Rather than optimizing individual mission tasks independently, our design philosophy focuses on integrating navigation, aerial observation, target localization, and payload delivery into a single coordinated autonomous system. The platform is designed to rapidly survey disaster-affected areas, generate accurate situational awareness, and assist rescue teams by providing precise target localization and autonomous delivery of emergency supplies. Achieving this objective requires not only accurate autonomous navigation but also seamless interaction between the mechanical, electrical, avionics, and imaging subsystems throughout the entire mission. Reliability was established as the primary design driver from the earliest stages of development. The aircraft incorporates a modular mechanical architecture, a navigation system based on the Global Navigation Satellite System (GNSS), stabilized aerial imaging, and a single digital link carrying both telemetry and video to ensure dependable operation under demanding environmental conditions. Special attention was given to structural integrity and maintainability in order to reduce operational risk and enable rapid field servicing when required.

Our engineering strategy emphasizes system-level integration rather than isolated subsystem optimization. Every design decision, from component placement and power distribution to payload deployment and communication architecture, was evaluated according to its contribution to overall mission success, reliability, and ease of operation. This holistic systems engineering approach enables the UAV to perform autonomous navigation, high-resolution aerial observation, accurate target localization, and controlled payload delivery as a unified and dependable mission platform. Ultimately, our objective is not only to satisfy the competition requirements but also to demonstrate a practical autonomous UAV architecture that can be adapted for real-world disaster response scenarios where accuracy, reliability, and rapid deployment are of critical importance.

B. Task Integration Strategy

To maximize mission efficiency and ensure reliable autonomous operation, each mission is designed not only as a standalone function but also as an integrated part of the overall UAV architecture. Navigation, target observation, communication, and payload delivery are designed to work collaboratively throughout the mission, enabling information sharing between the onboard avionics systems and the Ground Control Station (GCS). This integrated approach enables accurate autonomous waypoint navigation, continuous aerial observation, real-time mission monitoring, and precise payload deployment under varying operational conditions.

Navigation data from the GNSS receiver is processed continuously by the flight control unit, allowing the aircraft to hold a stable autonomous trajectory while the remaining subsystems operate in parallel. Target observation runs concurrently: the gimbal absorbs aircraft motion so that image quality does not degrade during maneuvers, and the same digital link carries both the telemetry stream and the live video the operator uses for mission monitoring and target verification. Because detection is performed onboard rather than at the ground station, confirmed target coordinates reach the guidance logic immediately and do not depend on link quality or operator reaction time.

Payload release is commanded from the same onboard sequence, which is what allows a target to be located and serviced within a single pass. The two delivery mechanisms are independent, so the mannequin and tent targets can be serviced at different stages of the flight without reconfiguring the aircraft between them; the mechanism is described in Section II-D. Treating navigation, observation, communication, and payload delivery as one coordinated workflow rather than as separate operations is what keeps mission complexity manageable for the operator, who supervises the mission instead of flying it.

C. Task Selection and Scoring Strategy

We attempt all five mission demonstration tasks in a single autonomous flight, and the vehicle was sized against their scoring functions rather than against generic performance targets.

The 6.7 kg (14.8 lb) all-up weight with both delivery objects installed falls below the 15 lb threshold that earns maximum Design for Rapid Response weight points, against the 35 lb limit. Each of the six 2S 13 Ah packs stores 96.2 Wh at its 7.4 V nominal voltage, below the 100 Wh per-battery limit; this is why the 6S2P bus is assembled from six small packs rather than fewer large ones, at no cost in total energy. The airframe measures $21.2 \times 21.2 \times 11$ inches in flight configuration and folds to $13 \times 13 \times 7.1$ in, which fits inside the $18 \times 14 \times 8$ in Personal Item sizing tier, and it is unpacked from that state to a flight-ready configuration in approximately three minutes using hands only.

The mission is flown by two operators, a Safety Pilot and a GCS Operator, which earns the full Efficient Operators score. This crew size is possible because detection and terminal guidance run onboard: the operator supervises the mission rather than flying it. Takeoff and landing are each commanded autonomously with a single command.

We plan eight autonomous waypoint laps before entering the Search Boundary, sized against a flight endurance of approximately 30 minutes at operational weight. Both delivery objects are carried simultaneously and both targets are serviced in a single pass, since landing to reload is not permitted, and the winch provides the retardant mechanism a scored delivery requires. Risk Mapping is served by FastMosaic, which produces the submission image on site within Mission Time.

Complexity versus reliability: Our available working hours were spent making a small number of capabilities dependable rather than adding capabilities we could not test. Three decisions follow from this trade. We rejected a P2 detection head and a dual-model ensemble because their computational cost was not repaid on a two-class detection problem. We kept tiled inference as a contingency rather than in the flight configuration, because it did not meet our onboard processing requirements in flight testing. Additionally, we selected a quadcopter over a hexacopter, accepting the loss of rotor redundancy in exchange for a mass budget that leaves room for payload and energy rather than for additional rotors. In each case a capability we could have added was traded for one we could verify before competition.

II. Design Strategy

Our design strategy was shaped by the operational requirements of disaster relief missions and the technical objectives of the SUAS competition. Each key mission requirement was translated into specific engineering objectives and addressed through an iterative design process emphasizing reliability, ease of maintenance, and system integration [1], [2]. The central trade was between payload capacity, flight duration, navigation accuracy, and structural integrity, all under the competition constraints; multiple design alternatives were evaluated and subsystem configurations were selected based on their contribution to overall mission performance. Subsystem compatibility, weight distribution, and thermal behavior were tracked throughout the development cycle alongside the functional requirements. The subsections below present the resulting configuration, propulsion, structural, airdrop, software, avionics, dataset, and mapping designs.

A. Configuration Selection and AHP Analysis

To rigorously determine the most suitable Unmanned Aerial Vehicle (UAV) architecture for autonomous post-disaster mapping and precision payload delivery, the Analytic Hierarchy Process (AHP) was utilized. Unlike conventional decision matrices, AHP mathematically synthesizes strict operational requirements to extract precise relative weights for competing architectures, namely the Quadcopter, Hexacopter, and Fixed-Wing configurations. The evaluation framework was constructed upon five primary criteria: hover stability, precision targeting, manufacturability, weight efficiency, and aerodynamic efficiency.

Given the strict operational constraints of the disaster response mission, which heavily relies on stationary precision to transmit exact victim coordinates and safely deploy payloads from a fixed altitude, hover stability and precision targeting were assigned the highest relative importance in the pairwise comparison matrix. Secondary factors, such as forward flight aerodynamic efficiency, were assigned lower weights, as high-speed cruising is subordinate to the absolute accuracy required during the critical payload delivery phase.

The AHP algorithmic output identified the Quadcopter as the optimal architecture, yielding the highest global priority score of 0.453 (Table I). Although the Hexacopter configuration (scoring 0.382) provided comparable hovering stability, its inherent mechanical redundancy introduced excess dead weight, which adversely affected the system's structural payload-to-weight ratio and overall manufacturability. Conversely, the Fixed-Wing architecture (scoring 0.165) was eliminated due to its intrinsic inability to meet the stationary hovering demands required for precision airdrops. Consequently, the Quadcopter was selected as it provides the most robust structural foundation, balancing mechanical simplicity with the lifting capacity necessary to safely manage a Maximum Takeoff Weight (MTOW) of up to 11 kg.

TABLE I
AHP GLOBAL PRIORITY SCORES FOR CONFIGURATION SELECTION

Alternative	Score	Rank	Justification summary
Quadcopter	0.453	1	Optimal balance of hover stability, robust structure, and payload-to-weight ratio
Hexacopter	0.382	2	Similar stability, but excess dead weight due to mechanical redundancy
Fixed-Wing	0.165	3	High aerodynamic efficiency, but intrinsically unable to meet stationary hover demands

B. Propulsion System

Propulsion was sized from the operating environment rather than from a catalogue. The competition site sustains winds of up to 13 knots (15 mph), and holding attitude and track against gusts of that magnitude while carrying both delivery objects set a thrust-to-weight ratio of 3:1 at operational weight as our design requirement. Meeting that ratio drove the selection toward a motor chosen for its efficiency rather than for peak thrust alone: a lower current draw at hover throttle extends the endurance on which the lap count of Section I-C depends. The four selected T-Motor U7 Lite V2.0 490KV units, sealed against dust and moisture for field operation, deliver a combined static thrust of 22.5 kg, giving 3.36:1 at the 6.7 kg operational weight and 2.05:1 at the 11 kg maximum takeoff weight.

The electronic speed controller was then specified above the current the motors draw at maximum thrust, so that transients at that operating point stay inside its rating rather than at its limit. A single RTF 80 A 4-in-1 unit was chosen over four discrete controllers because it removes three sets of housings, connectors, and wiring from the mass budget that Section II-C protects.

Six Leopard Power 13,000 mAh 25C 2S LiPo batteries coupled in a 6S2P arrangement make up the power system. This arrangement stays within the maximum permitted takeoff weight while providing the necessary operating voltage and energy capacity, and the pack is operated within safe discharge limits.

C. Airframe

The airframe structure is designed to provide a stiff primary load path while maintaining low mass and high field serviceability. The primary structures consist of carbon-fiber tubes and G10 composite plates, which are connected by aluminum parts that form the motor mount and the foldable joints. G10 was selected for the plates because it provides high in-plane stiffness, excellent electrical insulation, and RF transparency. Unlike metallic materials, G10 has minimal impact on radio-frequency signals, allowing GNSS, telemetry, and RC communication signals to pass through the structure with negligible attenuation. This makes it particularly suitable for mounting antennas and wireless communication components inside the airframe while maintaining reliable signal quality. The main purpose of the plates is to house the internal components, while the bottom side carries the battery pack and the payload. The foldable landing gear is attached to the outer ends of the carbon-fiber tubes, immediately adjacent to the motor mounts. Spacers are installed between the main plates to increase the bending stiffness of the UAV body and maintain structural rigidity under loading. Secondary structures like 3D-printed brackets and mounts are used for integration and protection. Thrust loads are transferred from the motor mounts through the carbon-fiber tubes and foldable joints to the main plates, while landing loads are introduced through the landing-gear mounts located adjacent to the motor mounts and are carried through the carbon-fiber tubes to the main structure.

To ensure stable and balanced flight, we carefully manage the center of gravity (CG) relative to the four rotors. The battery, as the heaviest component, is centrally placed on the main plate directly beneath the thrust vector of the rotors. This positioning keeps the CG aligned with the force generated during thrust, supporting a well-balanced and predictable flight. In addition, other internal components, such as the avionics, payload mechanism, and structural elements, are arranged to keep the CG close to the aircraft's central axis. This intentional layout minimizes the need for large corrections and helps maintain smooth, stable flight.

Our design focuses on three key modular features to maximize field serviceability, enable tool-free assembly, and ensure that the drone can be compactly packed into a standard pilot bag. First, the carbon-fiber arms attach via foldable joints. These joints allow the arms to fold for transport and lock securely for flight. They connect to the main plate and can be replaced individually if damaged, without removing other structural elements. Second, the landing gear mechanism is inspired by a folding motorcycle side stand. Each leg is mounted on a pivot near the motor mount, allowing it to fold alongside the carbon-fiber arm for compact storage. Its telescopic structure provides adjustable leg length, while the spring-assisted mechanism helps maintain the landing gear in either the deployed or folded position. Third, the battery pack is secured using hook-and-loop fasteners rather than permanent fixtures. This ensures rapid battery replacement while holding the pack securely in place. Together, these three modular elements allow critical subsystems to be serviced or replaced efficiently, minimizing downtime and enabling tool-free assembly and disassembly; Fig. 1 shows the airframe in its deployed flight configuration and Fig. 2 the same structure folded for transport.

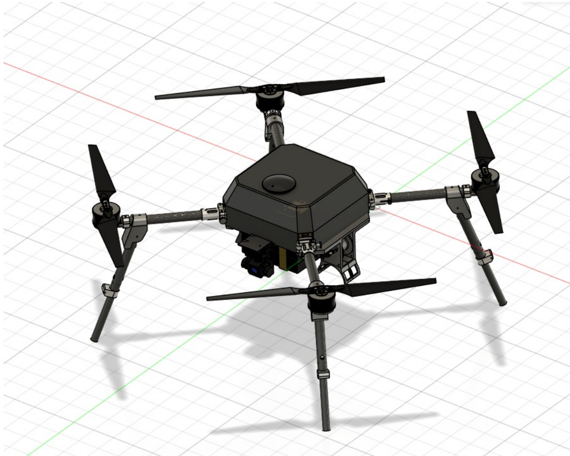

Fig. 1. CAD model of the UAV airframe showing the carbon-fiber arms, foldable joints, main plates, and the folding landing gear in the deployed configuration.

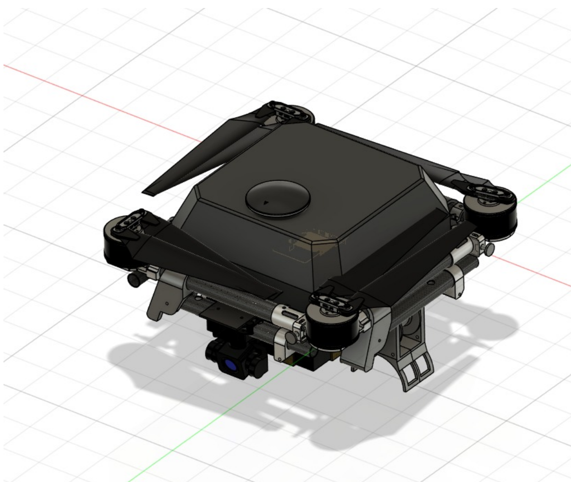

Fig. 2. The same airframe in its folded transport configuration, with the carbon-fiber arms retracted on their foldable joints and the landing gear telescoped in for compact packing into a standard pilot bag.

Structural simulation and fatigue assessment: To validate the structural integrity of the foldable arm and motor mounts under real operating loads, these components were analyzed in ANSYS under the thrust loading generated during flight (Figs. 3–6). Static analysis confirmed that both the foldable arm and the motor mounts maintain a safety factor of 1.5 against yield under maximum thrust load, providing margin for the repeated stress cycling introduced by continuous motor operation and the folding and unfolding cycles the joints undergo between flights. This simulation-based approach allowed us to validate the foldable arm and motor mount geometry before committing to final manufacturing, catching potential stress concentrations at the fold joints early in the design process rather than discovering them through field failures. Fatigue behavior under this cyclic thrust loading was also assessed to confirm that the foldable arm and motor mounts hold up over repeated competition use, supporting their durability alongside the portability they provide.

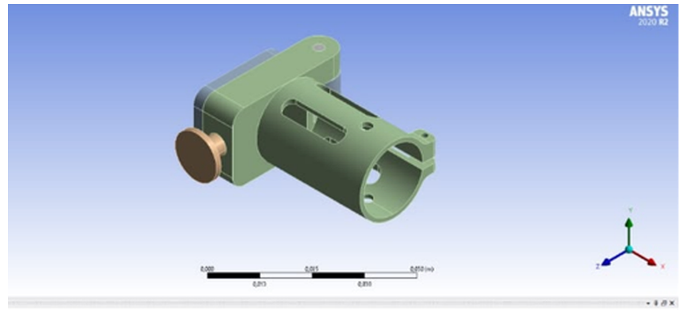

Fig. 3. Foldable arm and motor mount assembly as modelled in ANSYS, in wrought 6061-T6 aluminium.

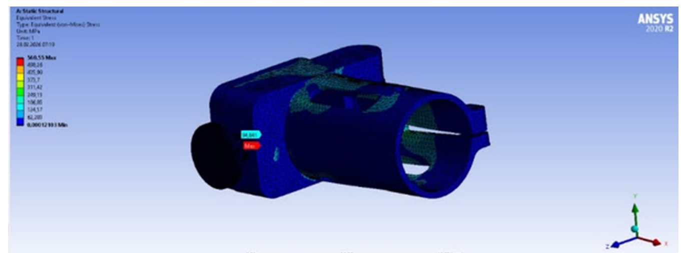

Fig. 4. Equivalent (von Mises) stress on the foldable arm and motor mount assembly under maximum thrust load.

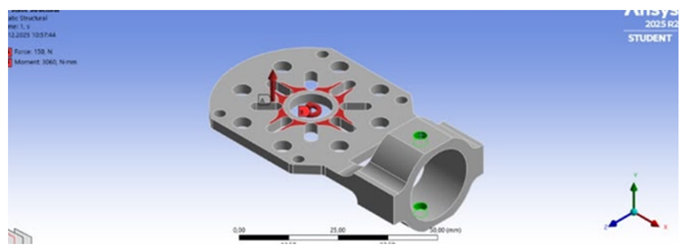

Fig. 5. Motor mount plate with the applied load case: a 150 N thrust force and the accompanying 3060 N·mm moment.

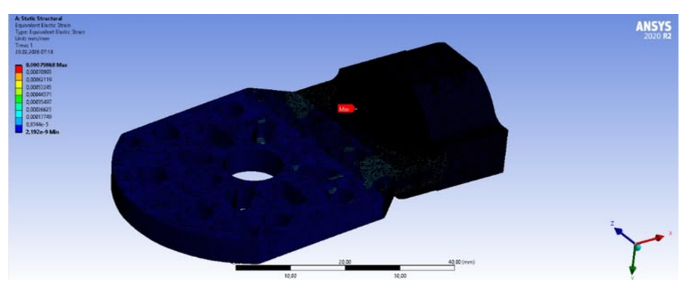

Fig. 6. Equivalent elastic strain on the motor mount plate under the same load case.

The structural and modularity choices above were made against a fixed mass budget, summarized in Table II. The airframe weighs 5.9 kg without payload and 6.7 kg with both payloads installed, against a maximum takeoff weight of 11 kg. That margin is what makes the battery pack sizing of Section II-B and the two-payload capacity of the airdrop module possible without approaching the structural limit, and it was also the deciding factor in selecting a quadcopter over a hexacopter in Section II-A, since the lighter configuration leaves more of the budget for payload and energy rather than for additional rotors.

TABLE II
VEHICLE MASS AND PERFORMANCE SUMMARY

Parameter	Value
Dry weight (no payload)	5.9 kg
Operational weight (both payloads)	6.7 kg
Maximum takeoff weight	11 kg
Maximum horizontal speed	18 m/s
Maximum vertical speed	5 m/s
Flight endurance (operational weight)	∼30 min
Maximum flight range	32 km

D. Airdrop

Our airdrop system uses a DC-motor-driven winch mechanism to enable independent payload deployment to the mannequin and tent targets, while a servo-actuated slider-crank mechanism securely locks the payloads in place until release. Each payload is secured by a dedicated spool lock actuated by a servo motor; while the lock is engaged, the payload remains fixed at the bottom of the UAV regardless of motor state. Once the UAV locks onto its target and positions itself directly above the designated drop location, the release sequence begins: the servo first opens the spool lock, and only then does the dedicated 800 RPM DC motor start, lowering the payload at a controlled, constant descent speed that eliminates free-fall. When the drop is completed, the motor stops. The entire release sequence is executed at the 50 m (164 ft) mission altitude, which keeps the aircraft above the 150 ft AGL minimum altitude fence throughout the delivery. Both motors are driven by a single HW-231 dual-channel H-bridge motor driver, allowing independent bidirectional control of each delivery mechanism. This configuration enables the two payloads to be deployed separately at different stages of the flight while providing reliable and precise control over the delivery process. The winch mechanism itself is built as a detachable module with a rail-guided sliding attachment. This allows the system to be quickly installed or removed without tools, while an integrated locking feature secures it firmly during operation.

E. Software Design

Our software architecture was developed to support autonomous mission execution while ensuring seamless integration between navigation, onboard perception, mapping, and payload delivery. Rather than treating these functions as independent software modules, the entire processing pipeline was designed around the operational requirements of the Search–Detect–Deliver mission, allowing multiple competition tasks to be completed during a single autonomous flight.

The software design began by translating mission requirements into measurable vision constraints. Since the SIYI A8 Mini camera has a fixed lens with digital zoom, target visibility is determined by flight altitude, zoom setting, and the resulting ground sampling distance (GSD). Instead of optimizing solely for inference speed, software requirements were derived from the minimum target pixel size necessary for reliable detection under realistic flight conditions. This approach guided the selection of image resolution, detection parameters, and processing architecture throughout the development process.

Ground sampling distance: With the gimbal locked at nadir, the ground footprint of a frame of width N_h pixels at altitude h above the target plane, and the resulting ground sampling distance, are

$$W_{\text{ground}} = 2h\tan\left(\frac{\text{HFOV}}{2}\right), \qquad \text{GSD} = \frac{W_{\text{ground}}}{N_h} \tag{1}$$

The mission is flown at 50 m (164 ft) AGL, which holds the aircraft above the competition's 150 ft AGL minimum altitude fence. With the 81° horizontal field of view of the SIYI A8 Mini at unity zoom and a horizontal frame width of 1920 px, (1) gives 44.5 mm/px at that altitude, which is too coarse for reliable detection of a supine mannequin. The camera's digital zoom recovers the resolution the altitude costs: the 4K sensor readout is downsampled by a factor of 2.1 into the 1080p stream at unity zoom, so cropping that readout at $3\times$ narrows the field of view to 31.8°, raises the effective focal length to $f = 3372$ px, and brings the ground sampling distance to 14.8 mm/px. At the selected zoom level, the mannequin occupies sufficient image area for reliable detection after resizing to the inference resolution. The target's short-axis footprint is the limiting factor and therefore determines the operational zoom setting at the mission altitude.

Object detection is performed onboard using a Jetson Orin NX companion computer running a YOLO11m-based detection pipeline. Live video from the SIYI A8 Mini camera is transmitted through an RTSP stream where only the most recent frame is processed to eliminate queue latency and maintain real-time performance. Confirmed detections are geolocated and transferred directly to the autonomous guidance system, allowing the aircraft to navigate toward the detected targets and initiate payload delivery without additional operator intervention.

To satisfy both detection accuracy and real-time processing requirements, several inference configurations were evaluated during development. While tiled inference improved small-object detection at longer distances, its computational cost caused onboard processing performance to fall short of what autonomous guidance required. Consequently, the operational flight configuration performs full-frame inference at 640-pixel resolution with a 0.25 confidence threshold. The 640 px input was set by the frame rate the guidance loop consumes rather than by accuracy alone, and it is the resolution against which the pixel budget above was sized, so the target footprint the detector receives is a designed quantity rather than a residue of downscaling. The 0.25 threshold is deliberately low because the release logic keys on persistence rather than on any single frame: arming requires the target to hold near the frame center across successive frames, so a detection that recurs at modest confidence is worth more than an isolated high-confidence hit, and the temporal requirement filters the false positives a low threshold admits. Alternative TensorRT-accelerated tiled inference remains available as a contingency configuration for recovering smaller or lower-confidence detections.

Target geolocation: Because the gimbal holds the camera at nadir and absorbs aircraft roll and pitch, the target pixel (u, v) in a $W \times H$ frame defines a ray that is already expressed in a locally level forward-right-down frame; it is brought into the local NED frame by the aircraft heading ψ alone, as reported by the flight controller,

$$\begin{aligned} \mathbf{r}_c = \Bigg[& -\frac{v - H/2}{H}\, 2\tan\!\left(\frac{\text{VFOV}}{2}\right), \\ & \frac{u - W/2}{W}\, 2\tan\!\left(\frac{\text{HFOV}}{2}\right),\ 1 \Bigg]^T, \\ \mathbf{r}_n = {} & \mathbf{R}(\psi)\, \mathbf{r}_c \end{aligned} \tag{2}$$

and intersected with the ground plane using the AGL altitude h, which yields the north and east offsets of the target and hence its coordinates. Here r_N, r_E and r_D are the north, east, and down components of $\mathbf{r}_n$, latitude and longitude are expressed in radians, and R_E is the local Earth radius:

$$\begin{aligned} [\Delta N,\ \Delta E]^T &= \frac{h}{r_D}\, [r_N,\ r_E]^T, \\ \text{lat}_t &= \text{lat} + \frac{\Delta N}{R_E}, \\ \text{lon}_t &= \text{lon} + \frac{\Delta E}{R_E \cos \text{lat}} \end{aligned} \tag{3}$$

Equations (1) and (2) are two projections of the same pinhole model: the per-pixel scale $2\tan(\text{HFOV}/2)/W$ in (2) is the inverse focal length $1/f$, with $f = (W/2)/\tan(\text{HFOV}/2)$ evaluated at the active zoom setting: 1124 px at unity zoom and 3372 px at the 3× setting at which detection and terminal guidance are flown. The ground sampling distance in (1) is that same factor projected onto the ground, $\text{GSD} = h/f$, and geolocation uses the zoom-matched value. Detection sizing and target geolocation therefore rest on one camera model rather than two independent calibrations.

Terminal guidance: A logged coordinate alone is not sufficient for release accuracy: it carries the geolocation error of a single detection, and the aircraft must be centered over the target at the instant of release. After navigating to the stored coordinate and settling, the aircraft therefore closes the loop on the live image rather than on the GPS fix, using image-based visual servoing. The controller drives the image-plane offset of the detected target from the frame center, $\mathbf{e} = (e_u, e_v) = (u - W/2,\ v - H/2)$, to zero. The offset is smoothed by an exponential moving average with $\alpha = 0.15$ to reject detection jitter, and is mapped onto the body axes of the nadir camera, in which image "up" is forward and image "right" is starboard. A discrete PID law running at 20 Hz ($\Delta t = 0.05$ s) then produces body-frame velocity commands, which are rotated into the NED frame by the current yaw ψ and streamed to the flight controller as MAVLink velocity setpoints:

$$\begin{aligned} \mathbf{v}_b[k] &= K_p\, \bar{\mathbf{e}}_b[k] + K_i\, \Delta t \sum_{j=0}^{k} \bar{\mathbf{e}}_b[j] \\ &\quad + K_d \left(\bar{\mathbf{e}}_b[k] - \bar{\mathbf{e}}_b[k-1] \right), \\ \bar{\mathbf{e}}_b &= [-\bar{e}_v,\ \bar{e}_u]^T, \\ [v_N,\ v_E]^T &= \mathbf{R}(\psi)\, \mathbf{v}_b \end{aligned} \tag{4}$$

The gains are $K_p = 0.003$ and $K_d = 0.004$ (m/s)/px with $K_i = 0.0001$ (m/s)/(px·s), where the derivative gain is defined per sample difference rather than per unit time, and the integral term is clamped to prevent wind-up while the target is briefly occluded. Command magnitude is saturated at 0.8 m/s so that the aircraft converges without overshooting. Detections are reported in the full 1920×1080 frame, so release is armed only after the target has remained within 20 px of the frame center, which by (1) corresponds to 0.30 m on the ground at the 50 m release altitude with the mission zoom setting; a target lost for a sustained period aborts the approach and returns the aircraft to the waypoint mission.

Control-loop limits: The controller uses a fixed-depth approximation of the classical interaction matrix, holding the estimated range at the release altitude rather than scaling the gain with measured range, so it is tuned for that single release altitude. The gains transfer to the mission altitude without rescheduling because the zoom setting, not the altitude alone, fixes the pixel-to-meter scale the loop acts on: at 14.8 mm/px the mission configuration sits within 17% of the 17.8 mm/px at which the loop was tuned. The releases reported in Section III-E were commanded from the stored geolocation alone.

F. Avionics and Electrical System

The avionics and electrical architecture were designed to ensure reliable subsystem integration while maintaining electromagnetic compatibility, thermal performance, and ease of maintenance. Power and signal cables were routed separately to minimize electromagnetic interference, and electronic components were arranged to improve cooling and accessibility. Each element was then selected against the requirement it had to satisfy rather than against a specification sheet.

An 11 kg maximum takeoff weight propulsion system produces high-frequency vibration that a single inertial sensor cannot reject, so the flight controller requirement was redundant, temperature-controlled inertial measurement with mechanical isolation; the Cube Orange Plus meets this with a triple-redundant heated IMU and runs ArduPilot natively.

Waypoint navigation and release positioning had to hold without a ground correction link, which set a standalone accuracy requirement on the GNSS receiver; the CUAV Neo3 Pro was chosen over the Here3+ on field reliability and cost at equivalent accuracy. Carrying telemetry and video on separate radios would have added a second receiver, its antennas and their wiring to a mass budget already committed to payload, so the SIYI HM30 was selected to carry both the MAVLink stream and HD video on one long-range link. The detection pipeline of Section II-E and the servoing loop that follows it exceed what the Jetson Nano and Xavier generations sustain, so the Orin NX was selected for the compute and memory those pipelines require within the power budget the airframe can carry. Both the mapping pipeline and the pixel budget of Section II-E assume the camera is held at nadir with aircraft motion absorbed, which made three-axis mechanical stabilization a requirement rather than a convenience; the SIYI A8 Mini provides it at a mass the payload budget can absorb.

In operation the Cube Orange Plus coordinates the navigation, propulsion, communication and payload subsystems, while the HM30 carries the A8 Mini video stream that the Jetson Orin NX processes onboard and that the operator monitors in Mission Planner.

G. Data Generation and Model Selection

The object detection system was developed using a dataset specifically curated to improve small object detection under realistic aerial operating conditions. A total of 16,015 images collected from multiple Roboflow projects were merged, duplicate images were removed using filename and perceptual hash comparison, and fifteen original classes were consolidated into two target categories: mannequin and tent. Vehicle-related classes were discarded while the significant class imbalance between mannequin and tent samples was reduced through oversampling of tent images. Training data were generated using overlapping 1280×1280 image tiles together with hard-negative background samples collected from debris and grass environments. Model validation was performed exclusively on real aerial imagery captured using a separate survey drone to ensure that performance metrics accurately reflected the operational environment rather than web-based datasets.

YOLO11m was selected as the object detection model because the primary limitation observed during initial experiments originated from dataset diversity rather than network capacity [3]. That finding set the model size: the smaller variants gave up accuracy we needed on the two-class problem, while the larger ones would have spent onboard throughput on capacity that was not the binding constraint, so the medium variant is the point at which further parameters stop paying for themselves. The model was trained using COCO-pretrained weights at an input resolution of 1280 pixels while data augmentation techniques including scaling, flipping, rotation, mosaic augmentation, and mixup were reintroduced to improve robustness against variations in object scale and viewing conditions. Alternative approaches, including a P2 detection head and a dual-model ensemble, were evaluated but ultimately rejected due to their increased computational cost without sufficient operational benefit for the two-class detection problem.

H. Risk Mapping

Our mapping subsystem converts the imagery collected during the autonomous waypoint tour into a single stitched map of the Search Boundary, generated on site immediately after the survey and written directly in the submission-ready image format (Fig. 7). The survey pass is flown at unity zoom, where the $81°$ field of view covers 85 m of ground per frame at the 50 m mission altitude and keeps the frame count of a full-boundary pass low; the $3\times$ setting of Section II-E is applied for target detection and terminal guidance. The design prioritizes fast on-site turnaround, robustness to the small frame count of a single survey pass (15–20 frames), and visual consistency across frame seams.

Fig. 7. Stitched map produced by the FastMosaic pipeline from survey imagery collected during a waypoint tour.

Pipeline (FastMosaic): Our in-house mosaicker, FastMosaic, converts the survey imagery into the deliverable map in four stages. GPS-assisted frame selection orders frames by the flight log and gates out frames taken during sharp maneuvers (heading change $> 30°$). Each remaining frame is matched against a sliding buffer of recent frames using ORB features [4] with CLAHE contrast normalization and placed with a 4-DOF similarity transform estimated by RANSAC,

$$\mathbf{x}' = s\,\mathbf{R}(\theta)\,\mathbf{x} + \mathbf{t} \tag{5}$$

Distance transform feathering, $w = d_{\text{new}}/(d_{\text{new}} + d_{\text{old}})$, blends each frame into the canvas, which is cropped and written directly in the .png format accepted for submission. GPS is used for selection and sequencing while placement is computed from visual features: geotags cannot provide pixel-level alignment, and features alone cannot provide ordering robustness; the design uses each for what it is good at.

We built our own mosaicker rather than adopting an external photogrammetry suite [5] because it keeps the mapping chain under end-to-end control inside the same Python toolchain as the rest of our software, adds no external processing dependency at the field station and produces the submission PNG directly. It also provides what batch photogrammetry tools do not: a real-time mode that consumes the A8 RTSP stream during the survey and builds an incremental preview, so a coverage gap costs a re-fly of one leg rather than a hole discovered too late in the submitted map.

Known limitations: FastMosaic output is a visual mosaic rather than a rigorously orthorectified product: similarity transforms do not model perspective or terrain relief, and incremental placement can accumulate drift over long chains, so we control quality at the source with nadir gimbal imagery, an overlap-rich survey pattern and the maneuver gate, validated by side-by-side comparison against reference imagery of the same area. Camera intrinsics use nominal datasheet-derived values ($f_x \approx 1124$ px) and lens distortion is left uncorrected. The similarity-transform model tolerates this at the overlap ratios we fly; in the geolocation chain of Section II-E it instead contributes to the error carried by a single logged coordinate, which is part of why release is designed to close on the live image rather than on the stored fix.

III. Testing Strategy

A. Overall Testing Strategy

Our testing strategy followed a structured and incremental verification process to validate the performance, reliability, and integration of the UAV before competition deployment. Rather than evaluating the aircraft only after full assembly, each subsystem was verified on its own, then in simulation, and only then in flight. This ordering reduced development risk, simplified troubleshooting, and ensured that no subsystem reached an autonomous flight test before satisfying its own operational requirements. The subsections below follow that order.

B. Component-Level Testing

Component-level testing was conducted to verify the functionality, reliability, and operational performance of each subsystem before full system integration. Each subsystem was evaluated independently under representative operating conditions to ensure compliance with mission requirements while minimizing integration risks during autonomous flight.

Propulsion System: The propulsion subsystem was tested to evaluate the performance of the motors, electronic speed controller (ESC), and power distribution system. ESC calibration was performed using Mission Planner in the ArduPilot environment to achieve consistent throttle response and synchronized motor operation. Battery voltage, current consumption, and power distribution were monitored under different operating conditions to evaluate the stability of the electrical system.

Flight Controller: The Cube Orange Plus flight controller was configured and calibrated using Mission Planner. Accelerometer, gyroscope, compass, and radio calibrations were completed before flight testing. Sensor outputs, flight modes, and failsafe functions were then verified during pre-flight inspections to ensure correct operation before autonomous flights.

Navigation System: The Neo3 Pro GNSS receiver was tested outdoors to evaluate its positioning accuracy and navigation performance. GPS signal acquisition and waypoint tracking performance were verified before autonomous flight operations.

Communication System: The HM30 digital transmission system was tested for telemetry communication, command transmission, and HD video streaming. MAVLink communication between the Ground Control Station and the flight controller was verified through Mission Planner by uploading test missions, monitoring telemetry data, and confirming stable communication throughout the tests.

Airdrop System: The payload delivery subsystem was tested by operating the servo-actuated locking mechanisms and the motor-driven winch assemblies independently. Lock release, controlled payload descent, and independent bidirectional control of each winch were verified for both delivery mechanisms, and repeated trials demonstrated consistent deployment.

Electrical Safety Considerations: The electrical system was designed to ensure safe and reliable operation under expected flight conditions. High-current propulsion wiring was routed separately from avionics power and signal lines to reduce electromagnetic interference. In addition, a hardware kill switch was incorporated to allow immediate shutdown of the propulsion system during emergency situations. Wire sizes were selected according to the expected current of each subsystem to minimize voltage drop and prevent excessive heating under load.

C. Simulation-Based Testing

Every new capability was flown in simulation before it was flown on hardware. The environment brings together Gazebo, ArduPilot SITL (Software-in-the-Loop), ROS 2 and pymavlink [6], and reproduces the actual flight architecture rather than running as a single desktop application: the flight dynamics simulator, the Mission Planner ground station and the Jetson companion computer operate as separate hosts linked by an Ethernet switch. Gazebo renders both the vehicle dynamics and the camera imagery, ArduPilot SITL flies the aircraft against those dynamics, Mission Planner commands the mission from the second host, and the Jetson consumes the rendered camera stream from the third, so the three exchange the same traffic they would in flight. Fig. 8 shows the simulation host during a waypoint run, where Gazebo, SITL, and the MAVProxy console share one machine and therefore its loopback interface; the ground station and the companion computer sit outside this view. Camera imagery reaches the Jetson over an RTSP stream, so the complete image-processing pipeline runs under conditions that closely match an actual flight. This approach cut development iteration time and limited the risk of damaging hardware during early testing.

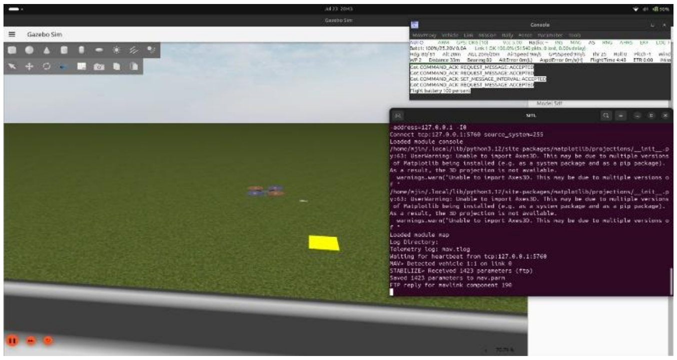

Fig. 8. The simulation host during a waypoint run, with the Gazebo world, the ArduPilot SITL instance, and the MAVProxy console sharing one machine. The Mission Planner ground station and the Jetson companion computer run on the two remaining hosts over the switch and are therefore outside this view.

Flight Path Simulation: SITL and Gazebo were used to simulate autonomous takeoff, waypoint navigation, landing, and Return-to-Launch (RTL) scenarios. Mission planning and monitoring were performed using Mission Planner, while survey missions were executed on waypoint grids generated with the Simple Grid tool. PID tuning and GPS-lock logic were validated in simulation using the actual UAV's weight and propulsion characteristics before hardware testing.

Object Detection Testing: Simulated objects with known positions and orientations were introduced into the Gazebo world so that detection accuracy, bounding-box alignment, processing latency, and the resulting onboard positioning decisions could be measured against ground truth on the target hardware.

Structural Simulation: The foldable arm and motor mounts were verified in ANSYS under maximum thrust loading before manufacture; the static and fatigue results are reported in Section II-C.

Emergency Scenario Simulation: Failure cases including GPS/EKF degradation, communication loss with the Ground Control Station or Jetson, and low-battery conditions were introduced using SITL. Automatic LOITER, landing, and Return-to-Launch behaviors were verified through telemetry logs, while the distributed network architecture enabled realistic testing of communication latency and link interruptions.

D. Dataset-Level Testing

The retrained YOLO11m model was evaluated using a validation dataset consisting exclusively of real aerial imagery captured using a separate survey drone. The dataset contained 31 images and 121 labeled instances. The model achieved an overall mAP50 of 0.950 and mAP50-95 of 0.745, while precision and recall were evaluated according to standard object detection metrics. Performance was assessed separately for mannequin and tent classes to verify detection reliability under realistic aerial conditions. Precision $P = TP/(TP + FP)$, recall $R = TP/(TP + FN)$; mAP50 averages AP at IoU 0.5 across classes (Table III).

TABLE III
VALIDATION PERFORMANCE OF THE RETRAINED YOLO11m MODEL BY CLASS

Class	P	R	mAP50	Val instances
mannequin	0.902	0.801	0.906	115
tent	0.942	1.000	0.995	6
All	n/a	n/a	**0.950**	121

Although the tent class achieved excellent detection performance, the evaluation included only six validation instances. Therefore, these results are reported transparently without claiming statistical confidence beyond the available sample size. In addition, tiled inference was investigated as a proof of concept for small-object recovery [7]. On two held-out aerial images where full-frame inference produced no detections, tiled inference recovered five and one detections respectively, at a peak confidence of 0.74. Onboard processing performance in this configuration fell short of what autonomous guidance required, which is why the full-frame configuration was selected for flight and tiling is retained only as a contingency.

E. Field Testing

Field validation proceeded in three stages. Detection robustness was first evaluated under low-light conditions using human stand-ins positioned on concrete surfaces. Across six trials, the detection system consistently produced confidence values of approximately 0.75 from distances of 20–25 m despite environmental conditions differing from the primary training dataset. Detection was then verified at the mission altitude: with the aircraft airborne at 49 m AGL and the camera at the 3$\times$ zoom setting of Section II-E, a human stand-in was detected with confidence values between 0.60 and 0.80, well clear of the 0.25 operational threshold. The reduction relative to the 20 m trials is the expected consequence of the smaller target footprint at altitude, and the retained margin is what allows the mission to be flown above the 150 ft AGL fence.

Finally, end-to-end flight tests were conducted to validate the complete autonomous mission pipeline under real operating conditions. The sequence was first exercised at approximately 20 m AGL to keep the release within visual range while the guidance chain was flown for the first time: both competition targets were detected, their positions were transferred to the guidance system, and autonomous navigation proceeded to the designated release locations, after which the beacon landed successfully on the tent target while the water bottle landed approximately 5 m from the mannequin. The same sequence was then repeated at the 50 m mission altitude with the 3$\times$ zoom setting, producing comparable detection confidence and a delivery offset of the same order. Both trials placed the payload well inside the 50 ft (15.2 m) radius that scores a delivery.

F. Mapping Testing

The mapping subsystem was verified through both component-level and mission-level evaluations. Prior to flight operations, RTSP video ingestion was validated using a dedicated bench-testing utility to ensure reliable communication with the SIYI A8 Mini camera. During map generation, frames captured during sharp heading changes were filtered automatically while images that failed feature matching were discarded instead of being forcefully integrated into the final mosaic, preserving overall map quality.

The generated stitched mosaic was evaluated according to the SUAS mapping requirements. Acceptance criteria included complete coverage of the surveyed area, seamless image alignment without visible stitching artifacts, consistent exposure across the mosaic, and visual comparison with reference imagery of the same region. These verification steps confirmed that the mapping pipeline produced competition-ready outputs suitable for autonomous mission planning and post-flight analysis.

REFERENCES

[1] M. H. Sadraey, *Aircraft Design: A Systems Engineering Approach*, ser. Aerospace Series. Chichester, U.K.: John Wiley & Sons, 2013.

[2] Q. Quan, *Introduction to Multicopter Design and Control*, 1st ed. Singapore: Springer, 2017.

[3] G. Jocher, J. Qiu, and A. Chaurasia, "Ultralytics YOLO11," https://github.com/ultralytics/ultralytics, 2024, accessed: 2026-07-26.

[4] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *2011 International Conference on Computer Vision*, 2011, pp. 2564–2571.

[5] OpenDroneMap Authors, "ODM: A command line toolkit to generate maps, point clouds, 3d models and DEMs from drone, balloon or kite images," https://github.com/OpenDroneMap/ODM, 2020, accessed: 2026-07-26.

[6] MAVLink Development Team, "pymavlink: A Python implementation of the MAVLink protocol," https://github.com/ArduPilot/pymavlink, 2024, accessed: 2026-07-26.

[7] F. C. Akyon, S. O. Altinuc, and A. Temizel, "Slicing aided hyper inference and fine-tuning for small object detection," in *2022 IEEE International Conference on Image Processing (ICIP)*, 2022, pp. 966–970.