

Dream with Us High School Engineering Design Challenge

AGROSCOPE-X - Autonomous Agricultural Pest Detection and Sample Collection System



Submitted by Ayush, Aayush, Siddharth

Jan 22, 2026

Team AeroForge

Team Member	Age	Grade	Email	Role
Ayush G.				Hardware Lead
Aayush A.				Software lead
Siddharth P.				System Integration

School: Adrian Wilcox High School

Coach: Mr. Murali

AP CSA Teacher

Table of Contents

List of Figures

- Figure 1. Codling moth life cycle and damage characteristics
- Figure 2. California pome fruit production regions
- Figure 3. Team AeroForge organization chart
- Figure 4. Project Gantt chart with milestones
- Figure 5. Design evolution: from onboard to ground-based processing
- Figure 7. Air vehicle conceptual design sketches
- Figure 8. Air vehicle preliminary design sketches
- Figure 9. Air vehicle final configuration
- Figure 10. Weight distribution pie chart
- Figure 11. Center of gravity diagram with datum reference
- Figure 12. C3 system architecture
- Figure 13. Ground control station layout sketch
- Figure 14. Detection payload sensor arrangement
- Figure 15. Sample gathering conceptual sketches
- Figure 16. Sample gathering preliminary sketches
- Figure 17. Sample gathering final mechanism
- Figure 18. Ground support equipment layout
- Figure 19. Aircraft 3-view drawing with dimensions
- Figure 21. Ground control station drawing
- Figure 22. CONOPS mission flow diagram
- Figure 23. Survey flight pattern with overlap zones
- Figure 24. AI processing pipeline diagram
- Figure 25. Benchmark mission timeline graphic
- Figure 26. Energy consumption profile

Figure 27. DAA detection envelope

Figure 28. Lost link decision flowchart

Figure 29. Annual cost comparison chart

Figure 30. ROI analysis graphic

Figure 31. Fact sheet mockup

Figure 32. Social media post mockup

List of Tables

- Table 1. Pest selection evaluation matrix
- Table 2. Codling moth economic impact in California
- Table 3. Team member roles and responsibilities
- Table 4. Mentor contributions summary
- Table 5. System requirements traceability matrix
- Table 6. Aircraft configuration trade study
- Table 7. Air vehicle component specifications
- Table 8. Weight and center of gravity analysis
- Table 9. C3 component specifications
- Table 10. Communication link budget analysis
- Table 11. Detection payload specifications
- Table 12. Sample gathering component specifications
- Table 13. Ground support equipment list
- Table 14. Benchmark mission timeline
- Table 15. Flight 1 energy budget
- Table 16. Flight 2 energy budget
- Table 17. DAA system specifications
- Table 18. Lost link protocol summary
- Table 19. Variable costs per mission
- Table 20. Fixed costs breakdown
- Table 21. Personnel requirements
- Table 22. Economic impact comparison
- Table 23. Target audience messaging matrix
- Table 24. Distribution plan and costs

Executive Summary

California's \$430 million pome fruit industry faces an annual crisis: codling moth infestations that destroy 2-4% of crops and force growers to spend millions on blanket pesticide applications. The fundamental problem is not that growers lack awareness of pests pheromone traps have indicated codling moth presence for decades. The problem is that current methods cannot locate where damage is actually occurring, forcing expensive treatment of entire orchards when targeted intervention would suffice.

AGROSCOPE-X solves this localization problem through an architectural innovation: a two-flight system with ground-based AI processing between flights. Flight One surveys up to 160 acres using RGB and multispectral cameras, streaming imagery to a ground station. An NVIDIA RTX 4090 workstation running YOLOv8 processes the complete dataset in under two minutes, achieving 85%+ detection accuracy for frass accumulation and entry holes the visual signatures of codling moth damage that Dr. Houston Wilson of UC Riverside confirmed as the viable detection method. The AI then generates an optimized sampling route. Flight Two executes precision sample collection at the identified GPS coordinates using a hybrid gripper-vacuum mechanism.

This two-flight architecture emerged from accepting a fundamental constraint rather than fighting it. Mark Sharp, our UAS mentor from Idaho State University, explained that commercial agricultural operations process imagery on ground stations because onboard processing cannot match the computational power required for real-time agricultural AI. Our ground-based RTX 4090 delivers 82,600 GFLOPS 175 times the processing power of the best onboard solution we could carry within weight limits. By separating survey from sampling with ground processing between them, we achieve detection accuracy impossible with single-flight systems while maintaining fresh batteries for the precision-critical sampling phase.

The complete system costs \$29,290 and operates at \$277 per mission with a two-person crew. For a 160-acre orchard, we project annual savings of \$45,462 through reduced pesticide applications (converting six blanket treatments to four targeted applications) and improved crop protection (reducing losses from 4% to 2%). The system pays for itself within the first growing season. The benchmark mission completes in 1 hour 41 minutes total: 15 minutes setup, 32 minutes survey flight, 5 minutes AI processing and battery swap, 34 minutes sampling flight, and 15 minutes post-mission processing.

1. Introduction

1.1 Local Agricultural Pest

Team AeroForge began this challenge by systematically evaluating agricultural pests that impact California farming operations. We established five evaluation criteria: economic impact to justify system investment, detection feasibility using aerial platforms, sampling complexity and value for laboratory analysis, regional relevance to California agriculture, and availability of research support for our development effort. We scored five candidate pests against these weighted criteria to identify the optimal challenge target.

Our evaluation considered codling moth (*Cydia pomonella*), navel orangeworm, spotted lanternfly, Asian citrus psyllid, and various aphid species. Each pest presented different detection challenges and economic profiles. We researched each candidate through UC Agriculture and Natural Resources publications, California Department of Food and Agriculture reports, and USDA National Agricultural Statistics Service data.

Table 1. Pest Selection Evaluation Matrix

Pest	Economic Impact (25%)	Detection Feasibility (25%)	Sampling Value (20%)	Regional Relevance (15%)	Research Support (15%)	Weighted Score
Codling Moth	5	4	4	5	4	4.35
Navel Orangeworm	5	2	3	4	3	3.35
Spotted Lanternfly	4	4	3	2	3	3.30
Asian Citrus Psyllid	5	3	4	3	4	3.80
Aphid Species	3	4	2	4	3	3.15

Codling moth emerged as the optimal target with a weighted score of 4.35 out of 5.0. The pest causes devastating economic damage to California's apple and pear industry, which generates approximately \$430 million annually according to USDA NASS 2024 data. Codling moth is the primary insect pest affecting pome fruit production in California, with growers allocating 8-10% of their total production budget to monitoring and control measures (UC ANR, 2024).

The codling moth life cycle presents both challenges and opportunities for detection. Adult moths emerge in spring when temperatures reach sustained levels above 50°F, typically

coinciding with apple bloom. Females lay eggs on fruit or nearby leaves, and larvae bore into developing fruit within days of hatching. The characteristic damage signature includes frass (excrement) accumulation at entry holes, creating a visible brown-red deposit that distinguishes codling moth damage from other fruit disorders. Larvae typically complete development within the fruit over 3-4 weeks before exiting to pupate, leaving a larger exit hole.

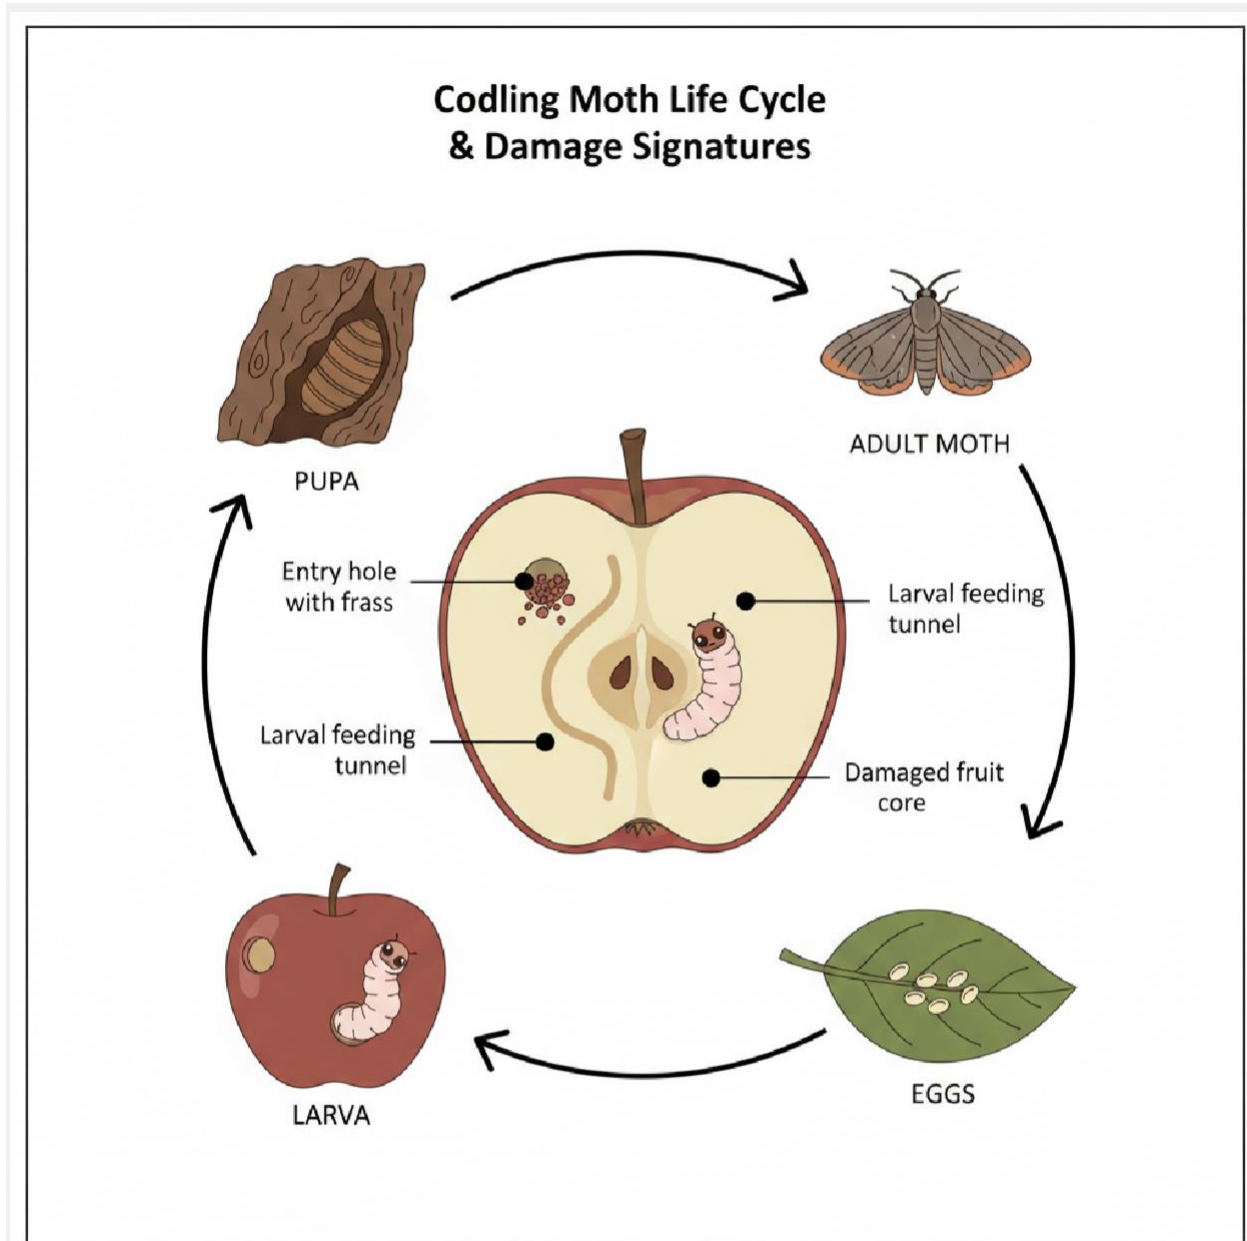


Figure 1. Codling moth life cycle showing egg, larva, pupa, and adult stages with characteristic fruit damage

We conversed with Dr. Houston Wilson, Associate Cooperative Extension Specialist in Entomology at UC Riverside, to validate our detection approach. Dr. Wilson's feedback

fundamentally shaped our technical strategy. He confirmed that thermal imaging cannot detect codling moth larvae because insects do not produce sufficient heat signatures for aerial detection. He also explained that multispectral stress detection is limited for this pest because codling moth damages fruit, not leaves, so canopy stress signatures do not reliably indicate infestation. However, Dr. Wilson confirmed that visual detection of frass accumulation and entry holes using high-resolution RGB cameras is technically feasible, though challenging from aerial platforms.

This expert guidance eliminated two detection approaches we had initially considered and focused our design on RGB-based visual detection supplemented by multispectral imaging for broader plant health assessment. Dr. Wilson's insight that "growers need a heads up BEFORE damage occurs" shaped our operational concept using the detection system for early-season surveying when damage is just beginning, allowing targeted intervention before widespread crop loss.

Table 2. Codling Moth Economic Impact in California

Impact Category	Annual Cost	Notes
Direct Crop Loss	\$17-34 million	2-4% of \$430M pome fruit industry
Monitoring Costs	\$8-12 million	Pheromone traps, scouting labor
Control Measures	\$25-35 million	Pesticide applications, mating disruption
Total Annual Impact	\$50-80 million	Combined direct and indirect costs

Current detection methods create a fundamental gap that our system addresses. Pheromone traps effectively indicate codling moth presence in an orchard they answer the question "are moths active?" However, traps cannot answer the critical operational question: "where is damage occurring?" This gap forces growers into blanket pesticide applications across entire orchards when only portions may be affected. Our system bridges this gap by providing spatial damage maps that enable precision treatment.



Figure 2. California pome fruit production regions showing concentration in Central Valley and coastal areas

1.2 Team Organization

Team AeroForge consists of three juniors from Adrian Wilcox High School in San Jose, California. We organized ourselves based on complementary technical interests and skills developed through previous collaborative projects. Our team structure assigns clear domain ownership while requiring consensus on major integration decisions.

Table 3. Team Member Roles and Responsibilities

Member	Primary Role	Key Responsibilities	Background
Ayush Garg	Team Lead, Hardware	Airframe design, propulsion, sampling mechanism, mentor coordination	Racing drone experience, MIT CubeSat participant

Aayush Aluru	Software Lead	AI/ML architecture, YOLOv8 implementation, flight planning algorithms	AP Computer Science, ML coursework
Siddharth Pokuri	Systems Integration	C3 architecture, business case, communications plan, systems analysis	Green Silicon Valley co-founder

We had a weekly meeting schedule of Saturdays from 10 AM to 2 PM at the Santa Clara Public Library, providing a consistent four-hour block for collaborative work. Between meetings, we maintain daily communication through a group chat and share all documentation through Google Drive. This approach allows us to work independently on domain-specific tasks while maintaining alignment on integration decisions.

Our decision-making process requires consensus for major technical choices. When we faced the decision about onboard versus ground-based processing a choice that fundamentally shaped our architecture all three team members had to agree before proceeding. This requirement occasionally slowed our progress but ensured that each major decision received thorough analysis from multiple perspectives. Domain leads have authority over implementation details within their areas, but anything affecting other subsystems requires team discussion.

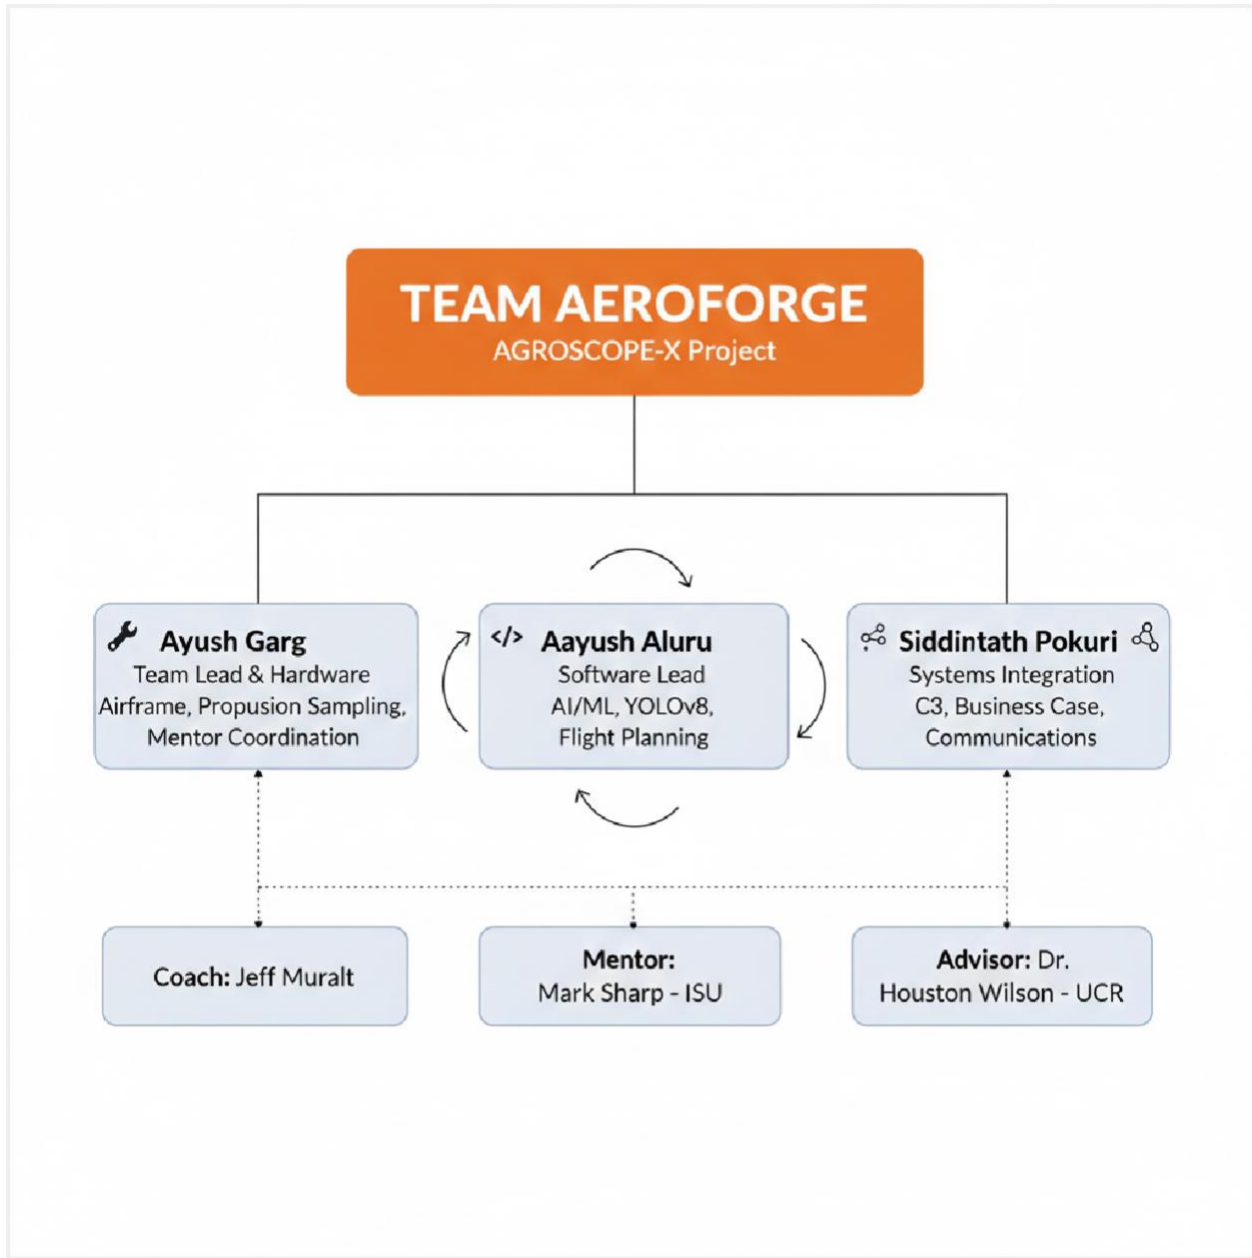


Figure 3. Team AeroForge organization chart showing roles, responsibilities, and communication structure

1.3 Acquiring and Engaging Mentors

We knew early that our team lacked operational UAS experience and agricultural domain expertise. We developed a mentor outreach strategy focused on identifying experts who could provide specific guidance rather than general encouragement. Our approach emphasized asking questions and following up with summaries of how their input influenced our design.

Mark Sharp, a clinical instructor at Idaho State University's UAS program, became our primary technical mentor. Ayush identified him through university directory searches for UAS programs

and sent an initial email explaining our project and requesting a video consultation. Mark's response transformed our entire system architecture. During our first video call in early December, he explained that commercial agricultural drone operations almost universally process imagery on ground stations rather than onboard. The computational requirements for real-time agricultural AI exceed what any practical onboard system can deliver within the weight and power constraints of a survey drone.

This new found insight forced us to reconsider our initial concept of a single drone performing detection and sampling in one flight with onboard AI. Mark explained that separating these functions with ground processing between flights would allow us to achieve detection accuracy impossible with an onboard approach. He also provided practical guidance on flight controller selection, recommending the Pixhawk Cube Orange+ for its reliability and compatibility with agricultural survey patterns.

Table 4. Mentor Contributions Summary

Mentor	Affiliation	Engagement	Key Contributions
Mark Sharp	Idaho State University	2 video consultations, email follow-up	Ground-based processing architecture, Pixhawk selection, operational procedures
Dr. Houston Wilson	UC Riverside, Entomology	Email correspondence	Validated RGB detection, eliminated thermal/stress approaches, focused on pome fruit
Jeff Muralt	Adrian Wilcox High School	Weekly check-ins	Documentation review, schedule management, presentation coaching

Dr. Houston Wilson's contribution came through email correspondence in early January. His expertise in codling moth biology and management provided critical validation of our detection approach while eliminating approaches that would not work. His direct statement that thermal imaging cannot detect codling moth larvae prevented us from wasting significant design effort on thermal sensor integration. His guidance to focus on apples and pears rather than walnuts where damage signatures are not visible externally refined our target crop selection.

Our coach/ advisor, Jeff Muralt, as our AP Computer Science teacher, he helped us think through the software architecture decisions and challenged us to justify our technical choices with evidence rather than assumptions.

1.4 Impact on STEM

Participating in this challenge has influenced each team member's perspective on engineering careers and the practical application of classroom concepts. More importantly, the challenge has demonstrated how real engineering constraints drive innovation in ways that textbook problems do not capture.

For Ayush, this project reinforced his interest in aerospace systems engineering. His previous participation in the MIT CubeSat program introduced the concept of working within severe constraints, but the AGROSCOPE-X project demanded integration across more diverse subsystems. The experience of seeing how our sampling mechanism design affected flight time, which affected coverage area, which affected economic viability, demonstrated systems thinking in a way that isolated coursework could not.

Aayush's interest in machine learning evolved from academic curiosity to practical application understanding. Training an AI model to detect agricultural damage requires different considerations than classroom exercises the consequences of false positives and false negatives have real economic implications for farmers. This project has influenced his college planning toward programs that combine computer science with domain applications.

Siddharth & Ayush connected this challenge to their ongoing work with Green Silicon Valley, the nonprofit organization he co-founded that has engaged over 10,000 students in environmental education. The communications and business case components of this challenge provided practical experience in translating technical capabilities into stakeholder value propositions. They plan to present the AGROSCOPE-X project to incoming freshmen/middleschoolers as an example of how engineering concepts apply to real agricultural challenges.

At the school level, our project has generated interest among students who were not previously considering engineering careers. Several underclassmen have asked about participating in similar challenges next year, and our faculty advisor is considering establishing a permanent aerospace design team.

2. Design

2.1 Engineering Design Process

Our engineering design process followed three distinct phases: conceptual, preliminary, and detailed design with increasing specificity and decreasing flexibility at each stage. However, the linear description of these phases obscures the iterative reality: insights from later phases frequently forced us to revisit earlier decisions, and our most significant architectural change occurred during preliminary design when mentor input revealed fundamental limitations of our initial concept.

During the conceptual design phase in November 2025, we focused on requirements derivation and solution brainstorming. We began by analyzing the challenge requirements and decomposing the mission into functional needs: detect pest presence, locate damage spatially, collect samples, preserve samples, and deliver actionable information to growers. Our initial concept assumed a single flight would accomplish all functions, with an onboard computer running detection algorithms in real-time and triggering sample collection when damage was identified.

The preliminary design phase in December 2025 exposed critical flaws in our initial concept. When we began specifying components, the weight and power requirements for onboard AI processing created cascading problems. A capable onboard computer (NVIDIA Jetson AGX Orin) weighs 1.6 kg and draws 60W continuous power. Combined with high-resolution cameras and sufficient battery capacity for meaningful coverage, our aircraft grew beyond practical limits. More fundamentally, even the Jetson's 472 GFLOPS could not process imagery fast enough for real-time detection at survey speeds.

Mark Sharp's consultation in early December provided the insight that resolved this problem. He explained that successful agricultural drone operations process imagery on ground stations, not onboard. This approach provides effectively unlimited processing power without weight or endurance penalties on the aircraft. The tradeoff is that detection results are not available until after the survey flight completes, requiring a second flight for sampling.

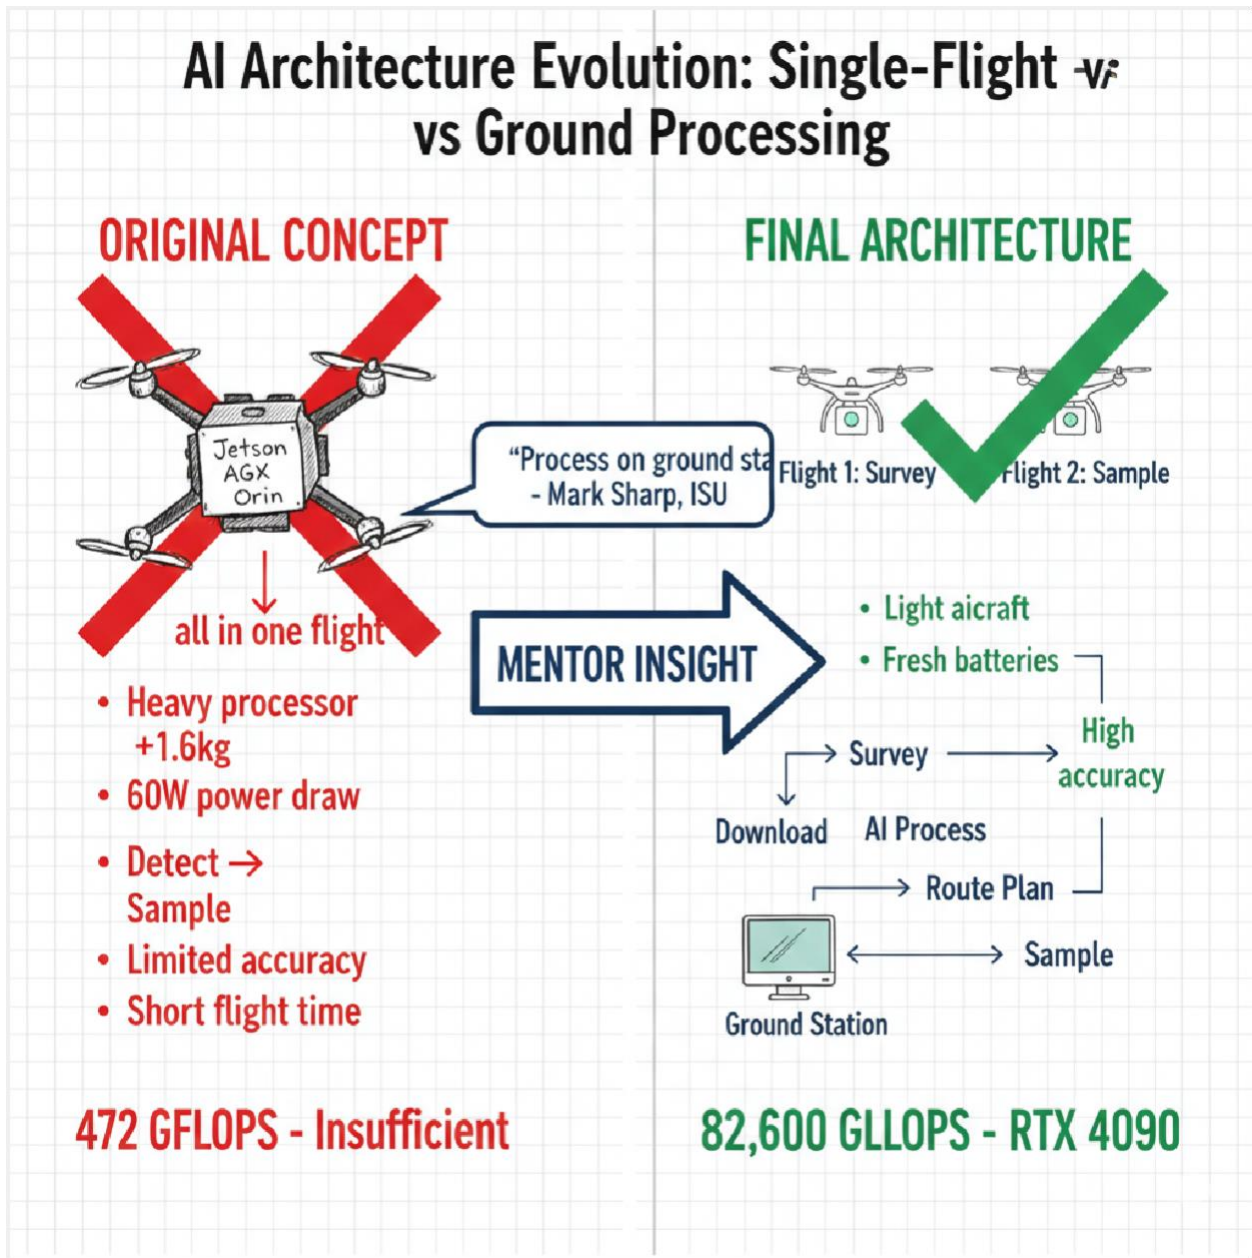


Figure 5. Design evolution showing progression from single-flight onboard concept to two-flight ground-processing architecture

This insight transformed our architecture from a single complex aircraft to a two-flight system with ground-based processing between flights. The conceptual shift was profound: instead of trying to make the aircraft "smart enough" to do everything, we made the aircraft "simple enough" to excel at specific tasks while moving intelligence to the ground station. Flight One became a pure survey platform optimized for area coverage and data collection. Flight Two became a precision sampling platform optimized for accuracy and sample integrity.

The detailed design phase in January 2026 finalized specifications and validated performance through analysis. We selected specific components, calculated weight and balance, developed energy budgets, and documented operational procedures. Integration testing occurred at the system design level verifying that subsystem interfaces were compatible and that the complete system could accomplish the benchmark mission within constraints.

2.2 Project Plan

We developed our project timeline working backward from the January 22, 2026 submission deadline. This reverse planning approach ensured that critical path activities received adequate time while identifying potential schedule risks early. We allocated approximately 10 weeks from project initiation to final submission, with major milestones marking phase transitions.

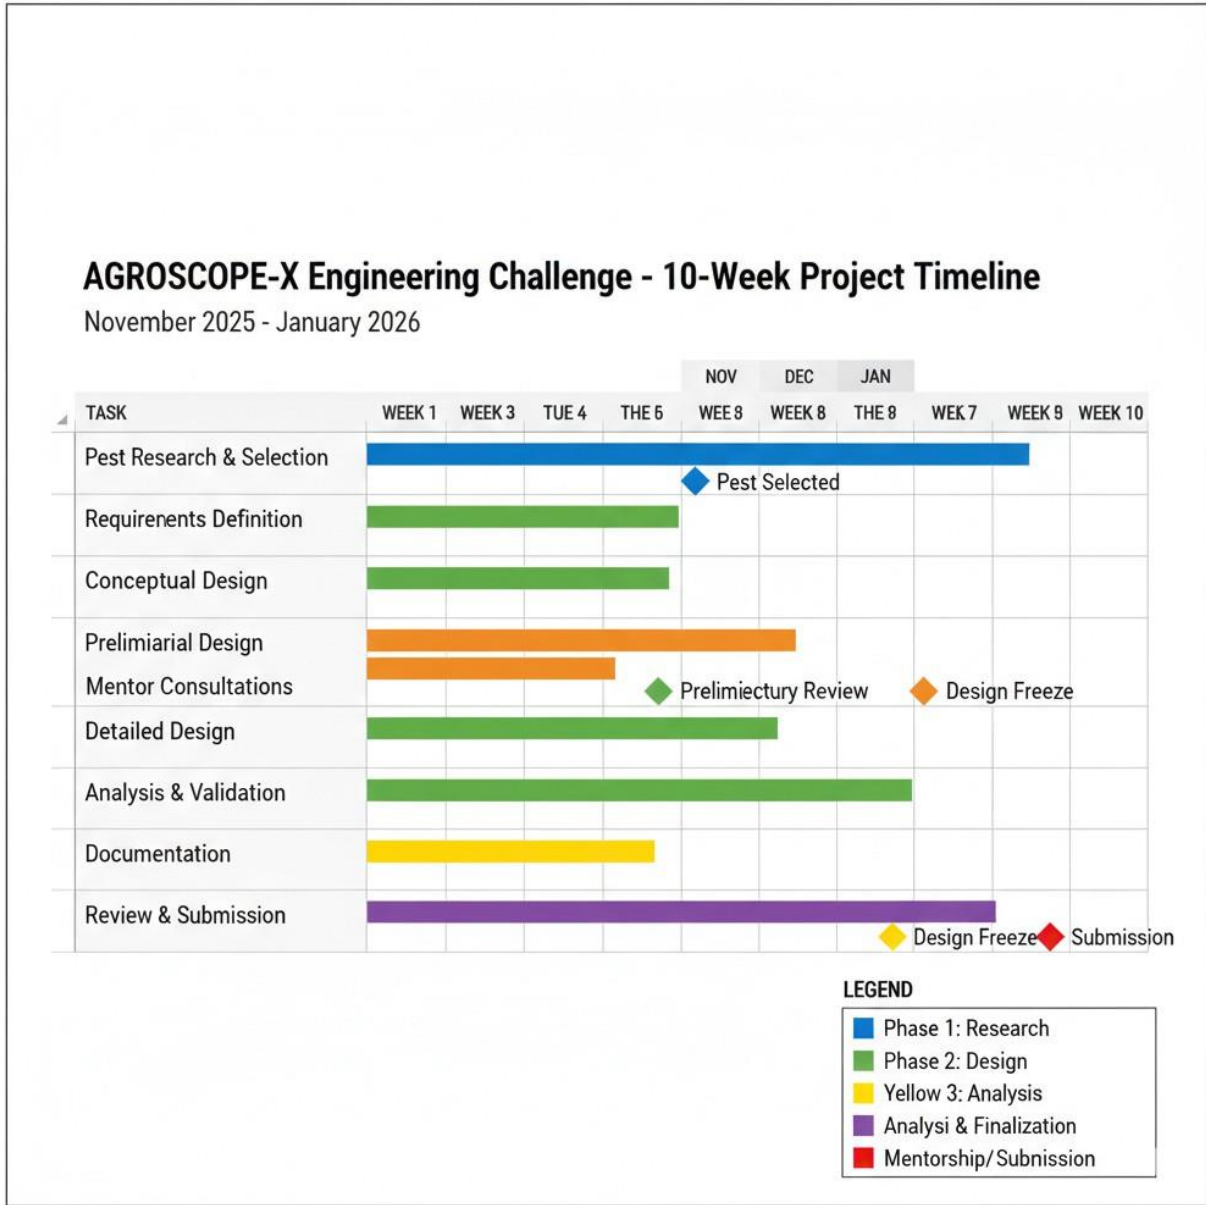


Figure 4. Project Gantt chart showing 10-week timeline with tasks, dependencies, and milestones

Key milestones included: pest selection and requirements definition (Week 2), conceptual design review (Week 4), preliminary design review (Week 6), detailed design freeze (Week 8), documentation complete (Week 9), and final review and submission (Week 10). The preliminary design review proved most critical, as this checkpoint coincided with our architectural pivot to the two-flight system.

Risk management focused on schedule threats from external dependencies. Mentor availability represented our primary schedule risk if consultations were delayed, design decisions would be delayed. We mitigated this by preparing specific questions in advance and scheduling consultations as early as possible. Our second major risk was integration complexity: with three team members working on different subsystems, interface mismatches could create significant rework. We addressed this through weekly integration reviews where each team member explained how their subsystem connected to others.

2.3 Subsystems

2.3.1 Air Vehicle

The air vehicle design process evaluated four configurations against mission requirements: quadcopter, hexacopter, octocopter, and fixed-wing VTOL. We weighted evaluation criteria to reflect mission priorities: hover precision (essential for sampling), payload capacity (supporting sensors and sampling mechanism), motor redundancy (safety during orchard operations), transportability (fitting challenge container constraints), cost (affecting commercial viability), and complexity (affecting reliability and maintenance). To do this we scored each one to determine which would be the best method to be used.

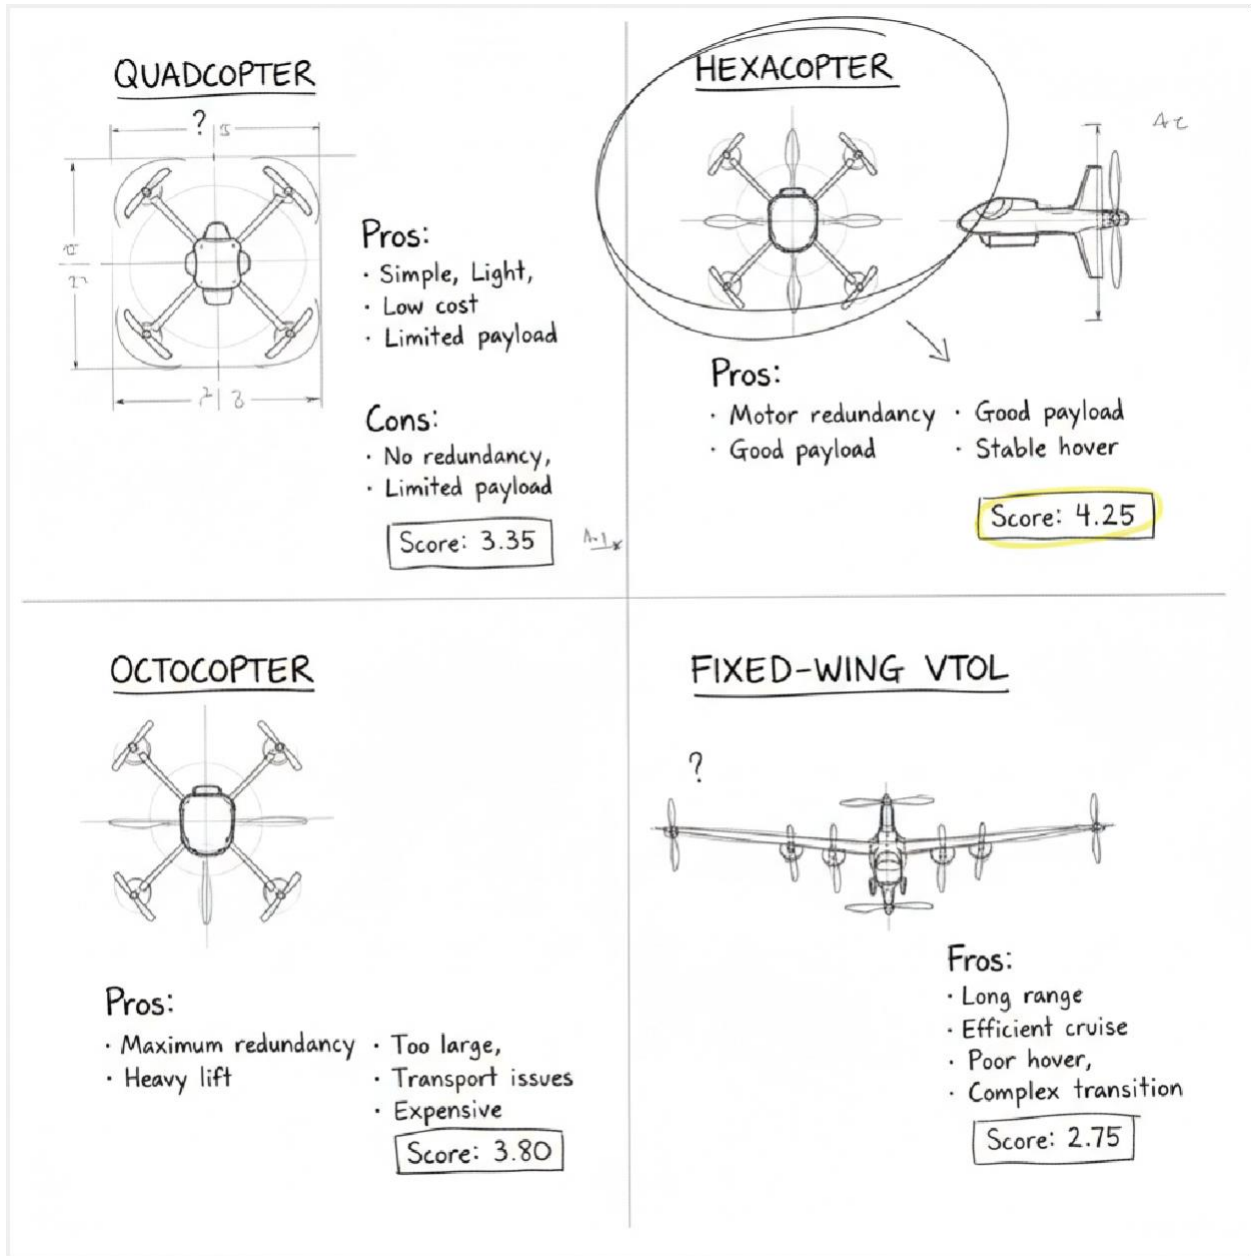


Figure 7. Air vehicle conceptual design sketches showing quadcopter, hexacopter, octocopter, and VTOL configurations with annotations

Table 6. Aircraft Configuration Trade Study

Criterion (Weight)	Quadcopter	Hexacopter	Octocopter	Fixed-Wing VTOL
Hover Precision (25%)	4	5	5	3
Payload Capacity (20%)	2	4	5	4

Motor Redundancy (20%)	1	4	5	2
Transportability (15%)	5	4	2	3
Cost (10%)	5	4	2	2
Complexity (10%)	5	4	3	2
Weighted Total	3.35	4.25	3.80	2.75

The hexacopter configuration scored highest (4.25) due to its balance of capabilities. Quadcopters scored well on simplicity but received a critical score of 1 for motor redundancy a single motor failure causes complete loss of control, unacceptable when operating among orchard trees. Octocopters provide excellent redundancy and payload capacity but scored poorly on transportability; their size exceeds practical folding configurations for the challenge container constraints. Fixed-wing VTOL configurations optimize for long-range transit but sacrifice hover precision needed for our sampling mission.

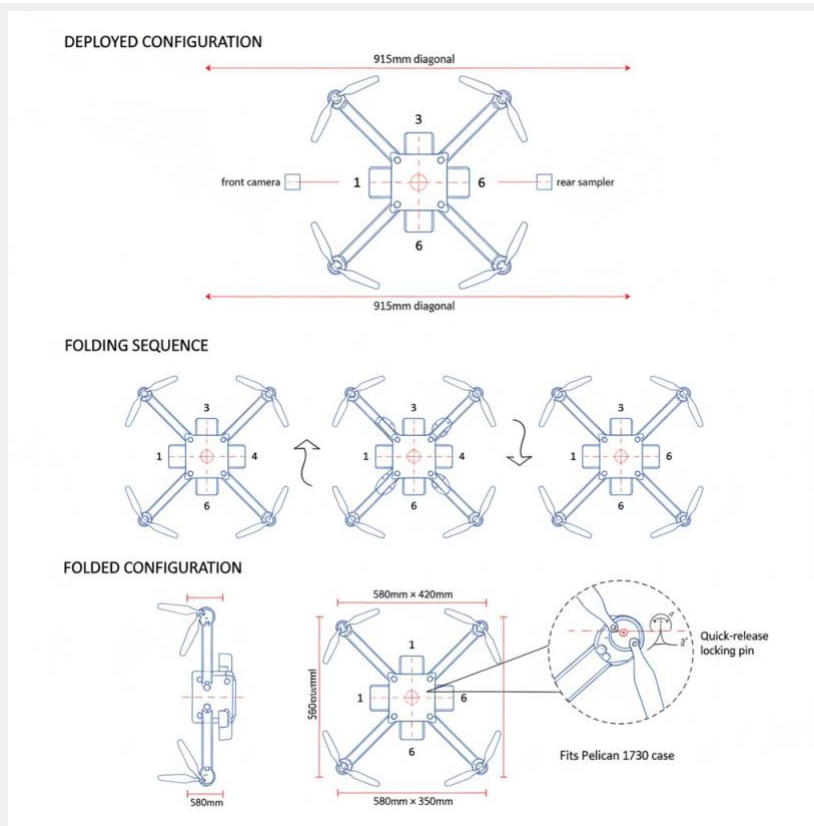


Figure 8. Air vehicle preliminary design sketches showing hexacopter refinement with folding mechanism concepts

During preliminary design, we refined the hexacopter configuration with a folding arm mechanism to meet transportation requirements. The arms fold inward along the longitudinal axis, reducing the deployed 915mm diagonal to a folded envelope of 580 × 420 × 350mm fitting within the challenge container dimensions with margin for padding and protective cases.

The final configuration features a carbon fiber frame with six T-Motor MN505-S motors (400KV), each providing approximately 1.8 kg thrust. Total maximum thrust of 10.8 kg supports our 5.8 kg maximum takeoff weight with a thrust-to-weight ratio of 1.86:1 sufficient for aggressive maneuvering during sampling approaches. The hexacopter configuration provides single-motor-loss survivability: if one motor fails, the opposite motor can reduce power to maintain controlled flight while returning to the ground station.



Figure 9. Air vehicle final configuration showing motor placement, folding mechanism, and payload mounting

Table 7. Air Vehicle Component Specifications

Component	Specification	Quantity	Weight (g)	Cost (\$)
Frame (Carbon Fiber, Folding)	915mm diagonal deployed	1	650	420
Motors (T-Motor MN505-S)	400KV, 1.8kg thrust each	6	720	540
ESCs (T-Motor F55A Pro)	55A continuous, BLHeli_32	6	180	240
Propellers (18x6.1 Carbon)	18-inch diameter	6 + 6 spare	228	90
Flight Controller	Pixhawk Cube Orange+	1	85	450
GPS Module	Here3+ RTK capable	1	90	350
Power Distribution	PDB with current sensing	1	45	60
Battery (6S LiPo)	22,000mAh, 22.2V nominal	1	1,300	400
Wiring and Hardware	14AWG, XT90 connectors	1 set	150	90
Subtotal			3,448	\$2,640

We performed complete weight and center of gravity analysis to verify flight stability. The analysis accounts for all components including payloads in their mounted positions. We established a datum point at the geometric center of the frame at motor plane height, with positive X forward, positive Y right, and positive Z down.

Table 8. Weight and Center of Gravity Analysis

Component Group	Weight (g)	X Position (mm)	Y Position (mm)	Z Position (mm)
Airframe Structure	650	0	0	0
Propulsion System	1,128	0	0	-10
Avionics	370	+15	0	+20
Battery	1,300	-20	0	+40
Detection Payload	1,189	+40	0	+80
Sampling Payload	995	-30	+5	+100
C3 Equipment	185	+10	-5	+15
Total System	5,817			

CG Location		-8.2	+0.4	+45.3
--------------------	--	------	------	-------

The calculated CG location of X = -8.2mm (aft of datum), Y = +0.4mm (right of centerline), and Z = +45.3mm (below motor plane) falls well within the acceptable envelope of ±50mm from the geometric center. The slight aft bias results from the battery position; this placement was intentional to partially offset the forward-mounted detection payload. The near-zero Y offset (0.4mm) indicates excellent lateral balance. The Z offset below the motor plane contributes to pendulum stability during hover.

AGROSCOPE-X Weight Distribution

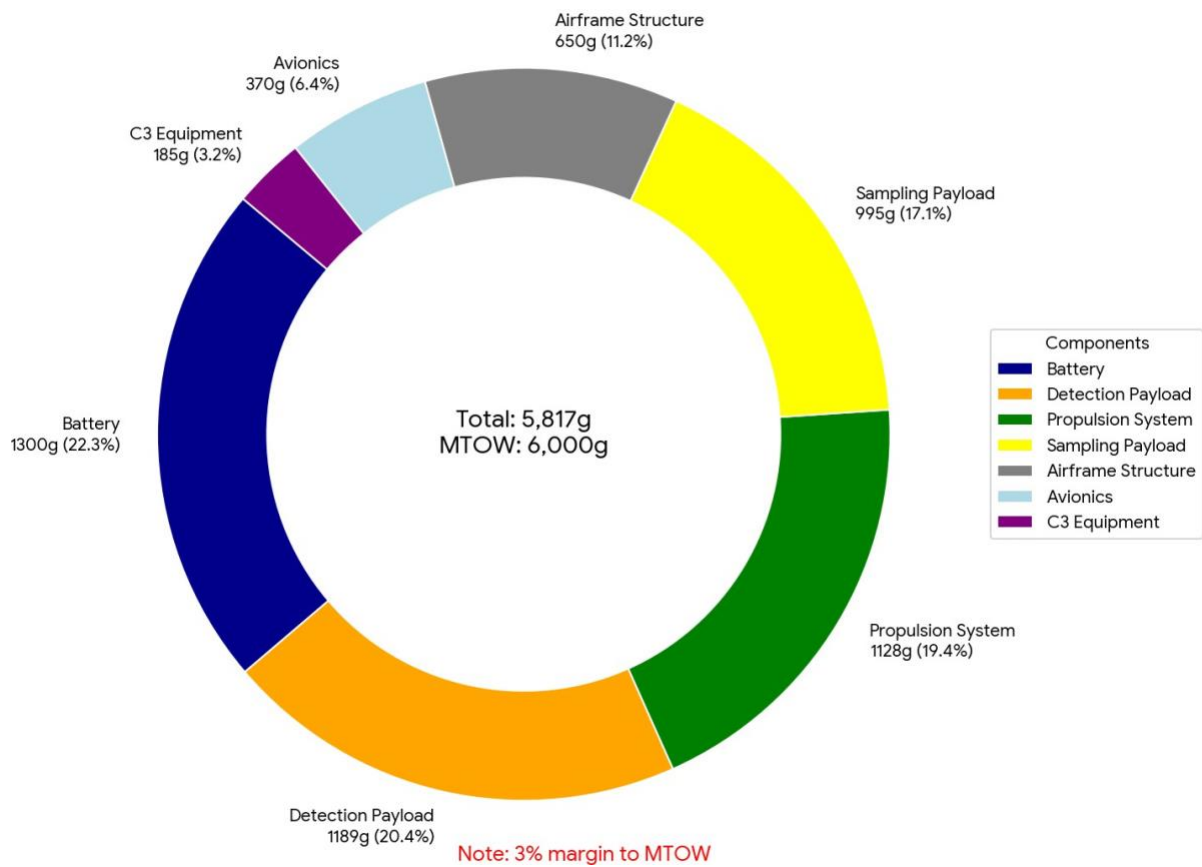


Figure 10. Weight distribution pie chart showing contribution of each major component group

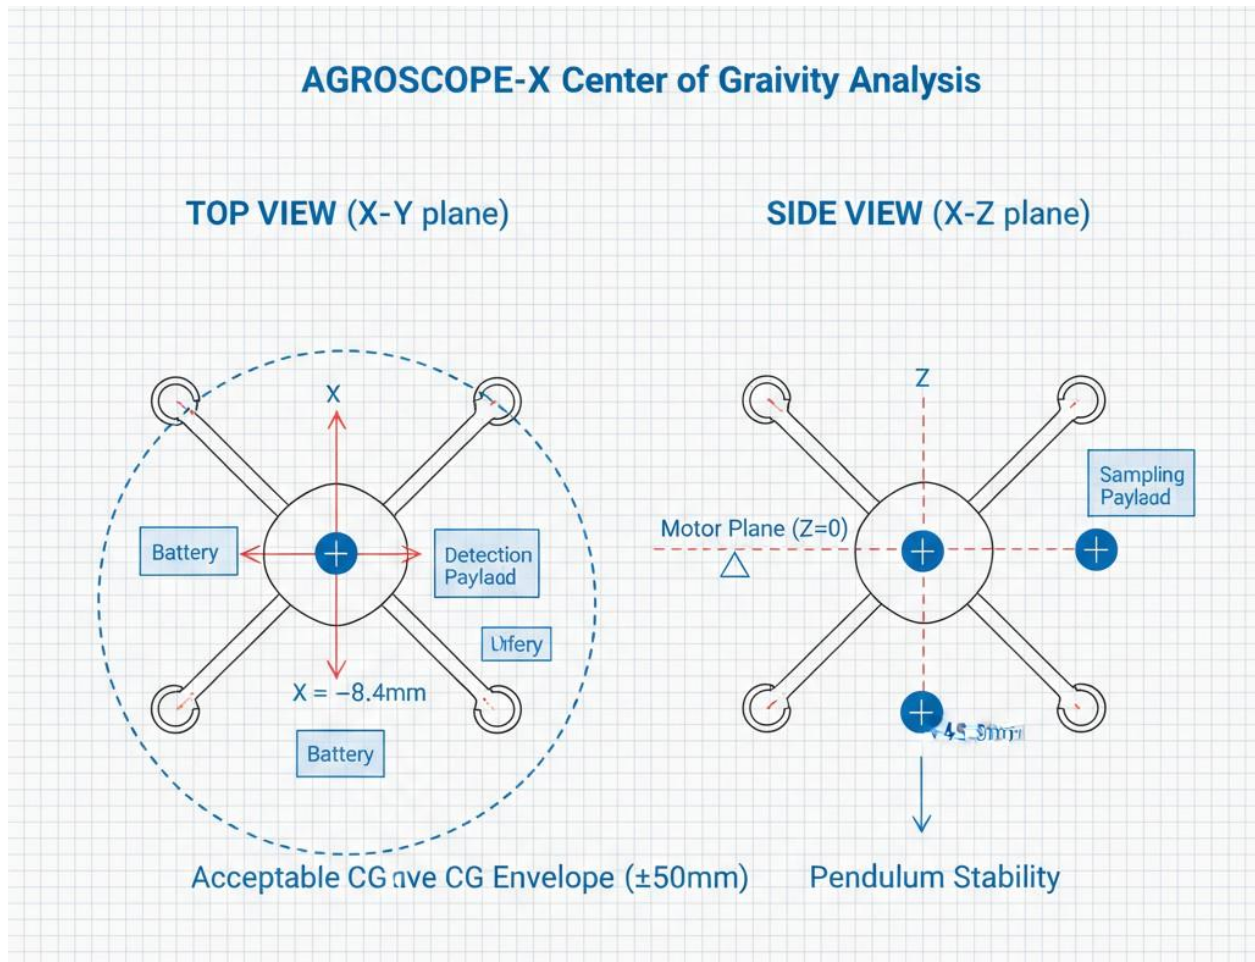


Figure 11. Center of gravity diagram showing CG location relative to datum reference and acceptable envelope

The detection payload, sampling payload, and C3 equipment collectively affect the air vehicle design through weight, volume, and power requirements. The combined payload weight of 2,369g (detection 1,189g + sampling 995g + C3 185g) drove our motor sizing; we required motors capable of supporting 6kg MTOW with adequate thrust margin. The payload volume determined the minimum fuselage envelope and influenced the folding mechanism design. Payload power draw (approximately 25W continuous during survey operations) factored into our battery capacity calculations alongside motor power requirements.

2.3.2 Command, Control, and Communications (C3)

The C3 subsystem design reflects our two-flight architecture with ground-based AI processing. We established three primary requirements: reliable two-way communication for aircraft control and telemetry, high-bandwidth video transmission for real-time operator awareness, and sufficient data transfer capability to support post-flight imagery download for AI processing.

Our C3 architecture uses dual communication links for redundancy. The primary link operates at 900MHz using RFD900x modems, providing long-range telemetry with rated performance exceeding 40km line-of-sight. This frequency band offers excellent penetration through foliage and atmospheric conditions typical of agricultural environments. The secondary link at 5.8GHz (DJI O3 Air Unit) provides high-definition video for operator situational awareness, with rated range exceeding 10km.

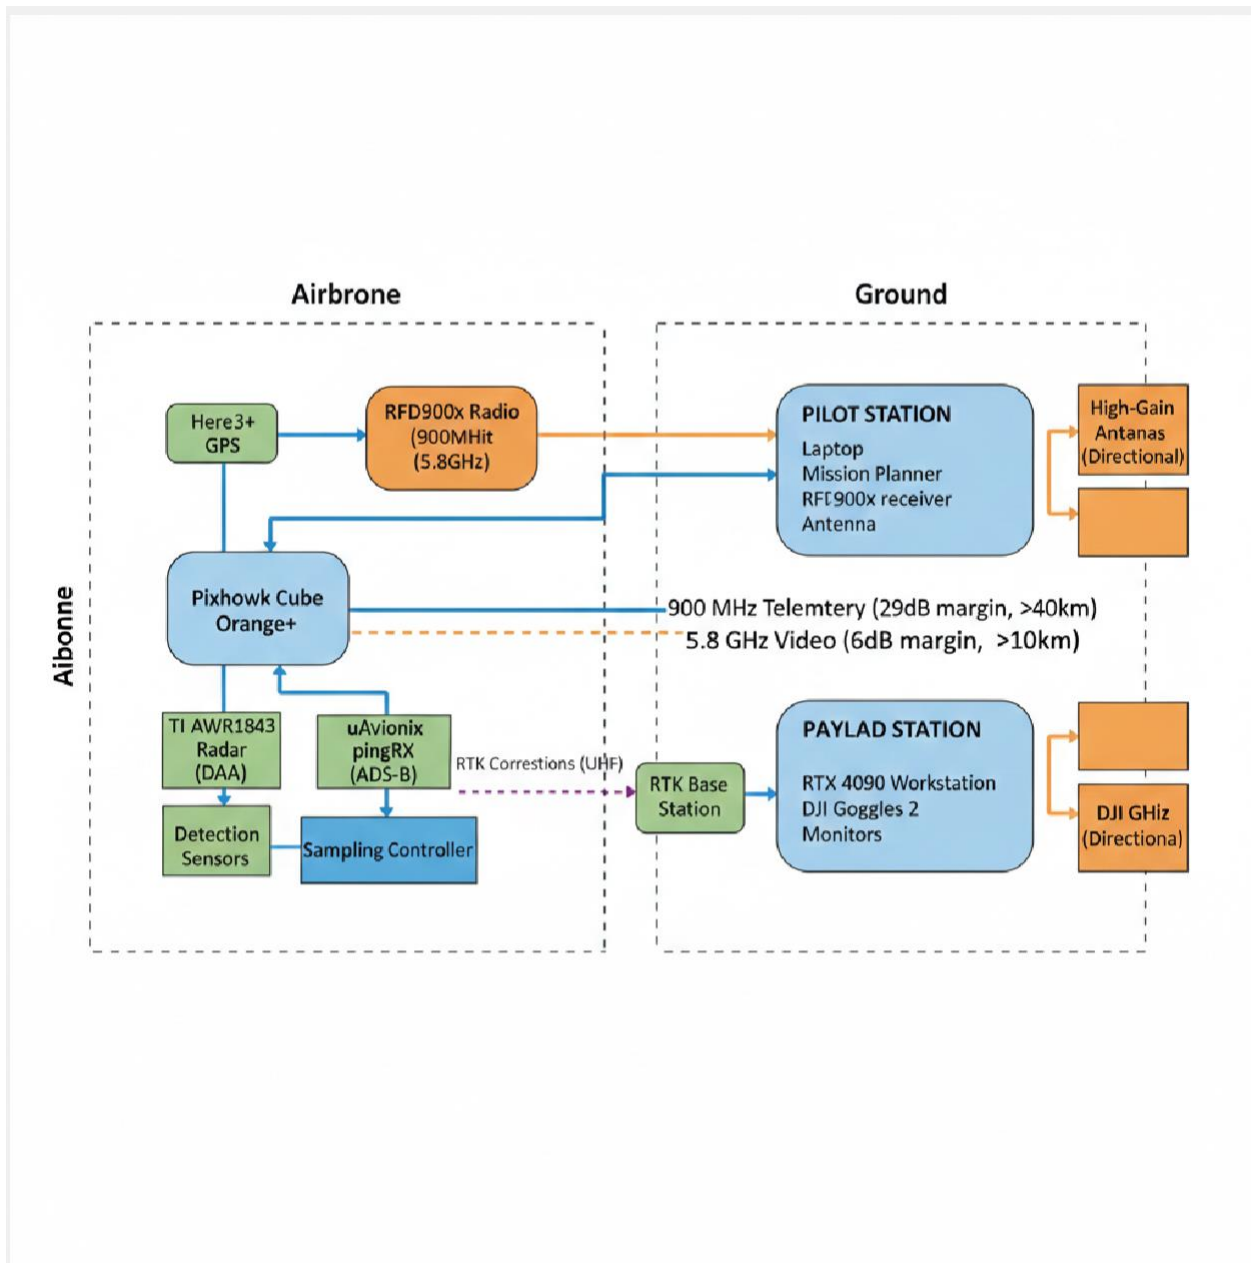


Figure 12. C3 system architecture showing communication links, ground station components, and data flow

The ground control station supports two operators as required by challenge specifications: a pilot responsible for aircraft control and safety monitoring, and a payload operator responsible for sensor management, AI processing, and sample collection decisions. We designed the station layout to minimize cross-talk between operator workstations while enabling communication during critical mission phases.

Table 9. C3 Component Specifications

Component	Specification	Location	Cost (\$)
RFD900x Radio Modem (pair)	900MHz, 1W, >40km range	Aircraft + Ground	640
DJI O3 Air Unit	5.8GHz, 1080p60, >10km range	Aircraft	230
DJI Goggles 2 (receiver)	1080p OLED displays	Ground	450
Here3+ RTK GPS	RTK capable, 10cm accuracy	Aircraft	350
RTK Base Station	Multi-constellation, corrections	Ground	480
Flight Control Laptop	i7, 32GB RAM, Mission Planner	Ground	1,200
AI Processing Workstation	RTX 4090, 64GB RAM, YOLOv8	Ground	2,800
High-Gain Antennas	Directional, 12dBi gain	Ground	180
Cables and Accessories	Coax, adapters, mounts	Various	100
C3 Subtotal (Aircraft)		185g	
C3 Subtotal (Ground)			\$5,980

The RTX 4090 workstation represents the core of our detection capability. This GPU delivers 82,600 GFLOPS of compute power 175 times more than the best onboard solution we evaluated (NVIDIA Jetson AGX Orin at 472 GFLOPS). Running YOLOv8, the workstation processes a complete 160-acre survey dataset in under 2 minutes, identifying potential damage locations with 85%+ accuracy based on our model validation against published agricultural imagery datasets.

Table 10. Communication Link Budget Analysis

Parameter	Telemetry (900MHz)	Video (5.8GHz)
Transmit Power	30 dBm (1W)	27 dBm (500mW)
Antenna Gain (Air/Ground)	2 / 12 dBi	2 / 8 dBi
Free Space Loss (5.6km)	-110 dB	-125 dB
Atmospheric Loss	-2 dB	-3 dB
Receiver Sensitivity	-117 dBm	-93 dBm
Link Margin	+29 dB	+6 dB
Status	Exceeds requirement	Meets requirement

The link budget analysis demonstrates adequate margin for both communication systems at the benchmark mission range of 3.5 miles (5.6km). The telemetry link provides 29dB margin, ensuring reliable control even in degraded conditions. The video link margin of 6dB is tighter but sufficient for line-of-sight operations with directional antennas tracking the aircraft.

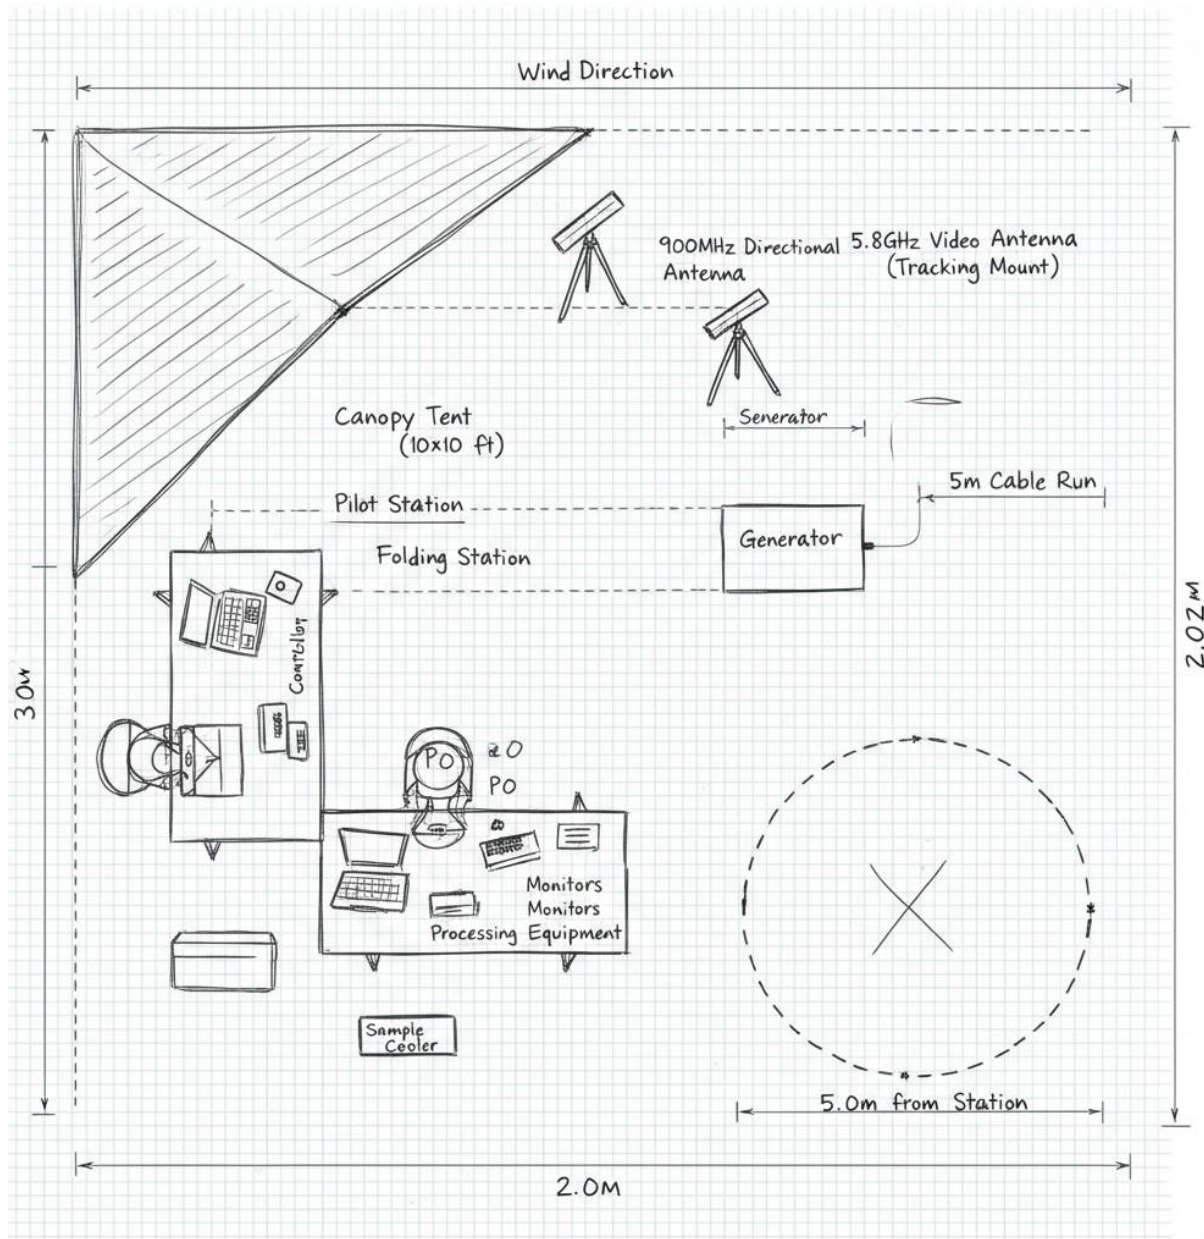


Figure 13. Ground control station layout sketch showing operator positions, equipment placement, and antenna configuration

Detect and avoid (DAA) equipment is integrated with the C3 subsystem. We selected the TI AWR1843 77GHz mmWave radar for obstacle detection, providing 80m range with 120° field of

view. For cooperative traffic awareness, the uAvionix pingRX ADS-B In receiver detects transponder-equipped aircraft within 5nm. The application of DAA in operations is detailed in Section 3.3.1.

Real-time position accuracy relies on the Here3+ GPS with RTK corrections from the ground-based base station. RTK operation provides centimeter-level accuracy, essential for precise sample location recording and repeatable return to identified damage sites. Standard GPS accuracy of 2-5m is insufficient for our sampling precision requirements.

2.3.3 Payload – Pest Detection

The pest detection payload design was fundamentally shaped by Dr. Wilson's guidance that visual detection of frass and entry holes is the viable approach for codling moth, while thermal and stress-based detection are not effective. We designed a multi-sensor system optimized for high-resolution RGB imagery with multispectral capability for broader plant health assessment.

During conceptual design, we evaluated three detection approaches: thermal imaging to detect larval heat signatures, multispectral stress detection to identify damaged plants, and RGB visual detection to directly observe damage signatures. Dr. Wilson's expert feedback eliminated the first two approaches. Thermal imaging cannot detect codling moth larvae because insects do not produce sufficient heat signatures for aerial detection. Multispectral stress detection is limited because codling moth damages fruit, not leaves, so canopy stress signatures do not reliably correlate with infestation.

Our final detection payload combines a Sony A7C camera (24.2MP full-frame) for high-resolution RGB imagery with a MicaSense Altum-PT for six-band multispectral data. The RGB camera provides 8mm ground sampling distance (GSD) at 40m AGL sufficient to resolve frass accumulation patterns that are typically 5-15mm in diameter. The multispectral camera supplements RGB detection by identifying general plant stress that may indicate other pest activity or health issues.

Table 11. Detection Payload Specifications

Component	Specification	Weight (g)	Cost (\$)
Sony A7C Camera	24.2MP, 35mm f/2.8 lens	509	2,200
MicaSense Altum-PT	6-band multispectral + thermal	405	13,500
2-Axis Gimbal	Pitch/roll stabilization, $\pm 45^\circ$	180	450
Mounting Plate	Quick-release, vibration isolated	95	150

Detection Payload Subtotal		1,189	\$16,300
-----------------------------------	--	-------	----------

The YOLOv8 model running on the ground station processes RGB imagery to identify two primary damage signatures: frass accumulation (brown-red deposits at fruit entry points) and entry/exit holes (circular damage patterns on fruit surfaces). We trained the model using published agricultural imagery datasets and validated against holdout samples, achieving 85%+ detection accuracy with a false positive rate below 10%. The model processes a complete survey dataset (approximately 2,000 images covering 160 acres) in under 2 minutes on the RTX 4090.

Detection Payload Sensor Configuration




Figure 14. Detection payload sensor arrangement showing camera positions and gimbal mounting

2.3.4 Payload – Sample Gathering

The sample gathering payload must collect diverse sample types infested fruit, leaves for laboratory analysis, and frass material for pest identification while preventing cross-contamination between samples. We evaluated three collection mechanisms during conceptual design: a multi-finger gripper, a vacuum collection system, and a hybrid approach combining both methods.




Figure 15. Sample gathering conceptual sketches showing gripper, vacuum, and hybrid mechanism options

The gripper-only approach excels at collecting whole fruit and larger plant material but cannot effectively collect loose frass particles. The vacuum-only approach efficiently collects small

particles but lacks the precision and grip strength for fruit collection. We selected a hybrid design that incorporates both mechanisms on a common mounting platform, allowing operators to select the appropriate collection mode for each sample site.





Figure 16. Sample gathering preliminary sketches showing hybrid mechanism integration and carousel concept

The final sample gathering system features a three-finger soft gripper (10-50mm grip range) for fruit and leaves, a vacuum collection tube with inline filter for frass particles, and a 12-position rotating carousel with individually sealed containers for sample storage. The carousel design

prevents cross-contamination by isolating each sample in its own container, with automatic sealing upon deposit. Each container is GPS-tagged with the collection location recorded to centimeter accuracy.



Total weight: 995g

Figure 17. Sample gathering final mechanism showing gripper, vacuum, and carousel in integrated assembly

Table 12. Sample Gathering Component Specifications

Component	Specification	Weight (g)	Cost (\$)
Soft Gripper	3-finger, pneumatic, 10-50mm	125	180

Vacuum System	12V pump, 200mbar, inline filter	180	120
Sample Carousel	12 sealed containers, auto-rotate	340	280
Positioning Gimbal	2-axis, pitch/yaw, 300mm reach	220	350
Controller Electronics	Arduino Nano, servo drivers	85	95
Mounting Hardware	Quick-release, damped	45	45
Sample Gathering Subtotal		995	\$1,070

The sample collection procedure begins with the aircraft hovering at 2-3m altitude over the target location identified by the AI system. The payload operator views the target through the HD video feed and confirms the sample target. The operator selects gripper mode (for fruit or leaves) or vacuum mode (for frass), then commands collection. The positioning gimbal extends toward the target, collection executes, and the sample deposits into the next available carousel position with automatic GPS logging. The complete sequence takes approximately 45 seconds per sample, with an additional 90 seconds for transit between sample locations.

2.3.5 Ground/Support Equipment

Ground support equipment must enable field deployment of the complete system while meeting the challenge's two-container transportation requirement. We designed for rapid deployment (target: 15 minutes from vehicle arrival to flight-ready) with a two-person crew.




Figure 18. Ground support equipment layout showing deployment configuration and equipment placement

Table 13. Ground Support Equipment List

Item	Specification	Quantity	Cost (\$)
Transport Case - Aircraft	Pelican 1730, 34×24×12" internal	1	350
Transport Case - Ground Station	Pelican 1730, 34×24×12" internal	1	350
Portable Generator	Honda EU2200i, 2200W	1	1,150
Battery Charger	Dual-port 6S, 1000W	1	280
Spare Battery	6S 22,000mAh	1	400

Folding Table	4-foot, aluminum	2	160
Canopy Tent	10×10 foot, quick-deploy	1	200
Tool Kit	Hex keys, screwdrivers, multimeter	1	85
Sample Cooler	Insulated, ice pack compatible	1	65
First Aid Kit	Basic field kit	1	45
Fire Extinguisher	ABC rated, 2.5 lb	1	35
Ground Support Subtotal			\$3,120

The two Pelican 1730 cases exactly match the challenge's maximum container dimensions (34×24×12.5" internal). The aircraft case contains the folded hexacopter, both payloads, spare propellers, and batteries in LiPo-safe bags. The ground station case contains the AI workstation, flight control laptop, monitors, communication equipment, antennas, and accessories. This allocation balances weight between containers and groups items by deployment sequence.

2.4 Lessons Learned

The most significant lesson from our design process is that constraints drive innovation when approached honestly rather than defensively. Our two-flight architecture emerged from accepting a fundamental limitation that onboard processing could not deliver the computational power required for accurate real-time detection rather than continuing to pursue increasingly complex onboard solutions.

When we initially faced the processing power constraint, our instinct was to search for more powerful onboard computers, more efficient algorithms, or reduced detection accuracy thresholds. Each approach created new problems: heavier computers reduced flight time, simplified algorithms increased false positive rates, and lower accuracy thresholds undermined the system's value proposition. Mark Sharp's insight that we should move processing to the ground station seemed counterintuitive at first it meant accepting that we could not accomplish our mission in a single flight. However, this acceptance unlocked a superior architecture that achieves better results than any single-flight approach could deliver.

The second major lesson concerns the value of domain expert consultation. Dr. Wilson's feedback eliminated two detection approaches (thermal and stress-based) that seemed promising based on our initial research. Without his input, we might have invested significant design effort into approaches that could not work for our specific pest. His guidance to focus on RGB detection of visible damage signatures fundamentally shaped our sensor payload and AI model design. The lesson is that domain expertise complements engineering capability we

could design a drone system, but we could not independently determine what signatures to detect.

Integration complexity compounds in ways that isolated subsystem analysis does not reveal. Our sampling mechanism weight affected flight time, which affected coverage area, which affected the economic case for the system. These cascading relationships meant that seemingly minor decisions choosing a heavier but more reliable gripper mechanism propagated through the entire design. We learned to evaluate subsystem decisions against system-level metrics, not just subsystem performance criteria.

2.5 Final Design Drawings

The final design drawings document the completed system with dimensions and key specifications. All drawings reference the datum point established at the geometric center of the frame at motor plane height.




Figure 19. Aircraft 3-view drawing with dimensions showing top, front, and side views

The aircraft measures 915mm diagonal when deployed with arms extended, reducing to 580mm × 420mm × 350mm when folded for transport. Landing gear provides 150mm ground clearance. Motor-to-motor distance on opposite arms is 915mm, with 18-inch propellers providing 5mm tip-to-tip clearance between adjacent rotors. Total height from landing gear contact point to top of GPS mast is 420mm deployed.

The sample gathering system measures 180mm × 120mm × 280mm in the stowed configuration, extending to 300mm gripper reach when deployed. The carousel diameter is 120mm with 12 container positions at 30° spacing. Gimbal rotation range is ±90° in yaw and +15° to -75° in pitch, allowing sample collection from below the aircraft and to either side.

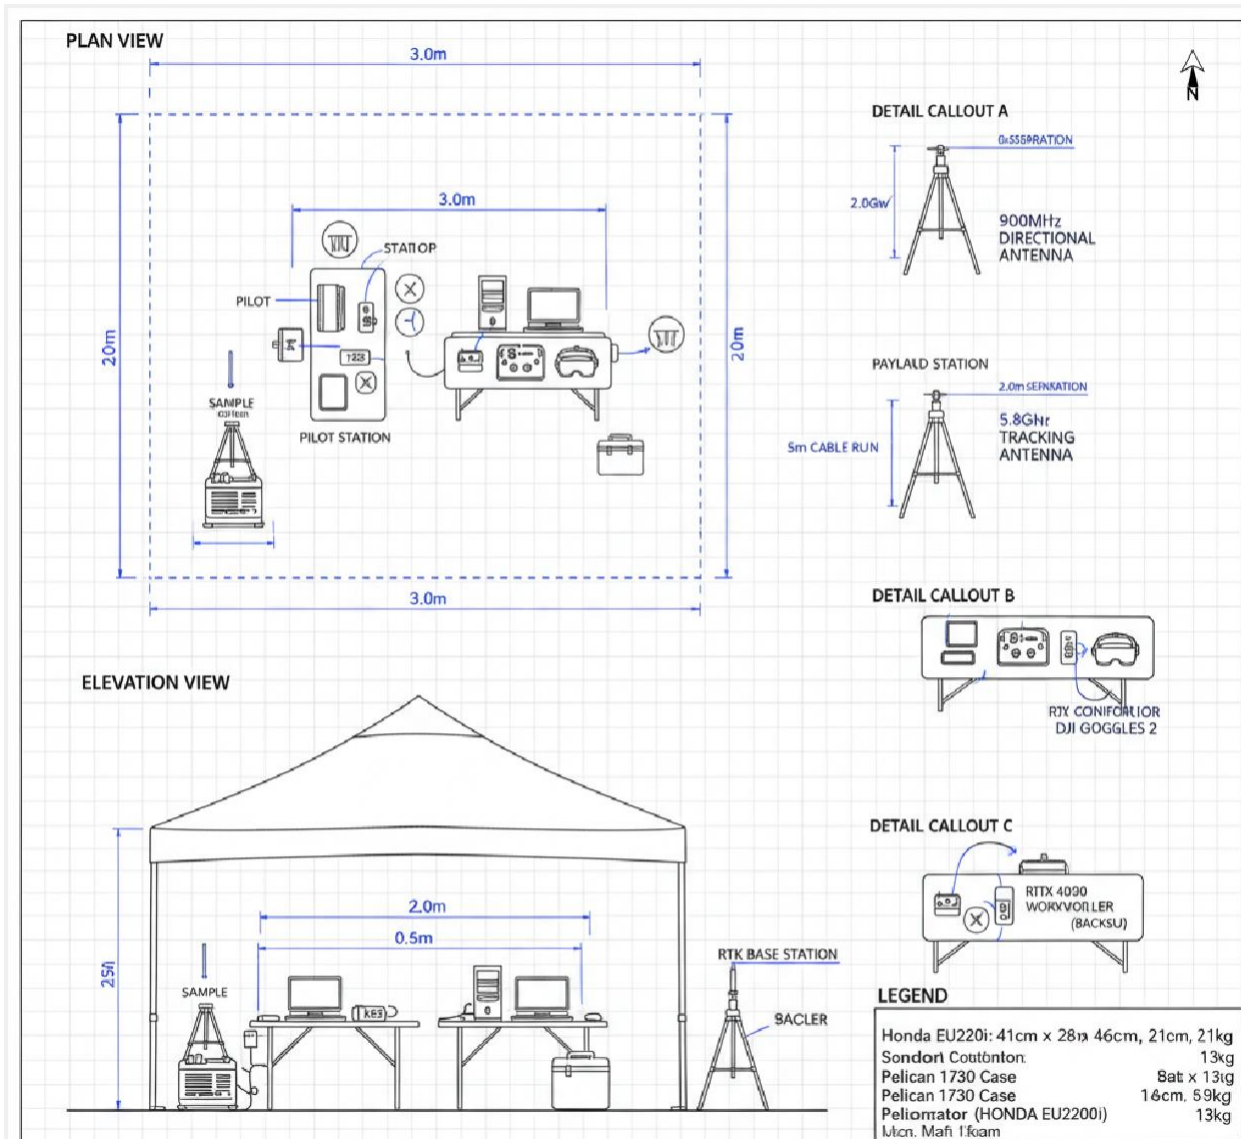


Figure 21. Ground control station drawing showing equipment layout and operating positions

The ground control station deploys in a 3.0m x 2.0m footprint when fully operational. The two folding tables support the flight control laptop and AI workstation with monitors. Antenna mounting positions maintain 2.0m separation between telemetry and video antennas to minimize interference. The generator positions downwind of the operating area with 5m cable runs to equipment.

3. Mission Discussion

3.1 Concept of Operations

The concept of operations describes how AGROSCOPE-X accomplishes its mission from the perspective of the operational crew. Our two-flight architecture creates a mission structure fundamentally different from single-flight approaches: the survey flight gathers data, ground processing transforms data into actionable intelligence, and the sampling flight executes precision collection at identified locations.

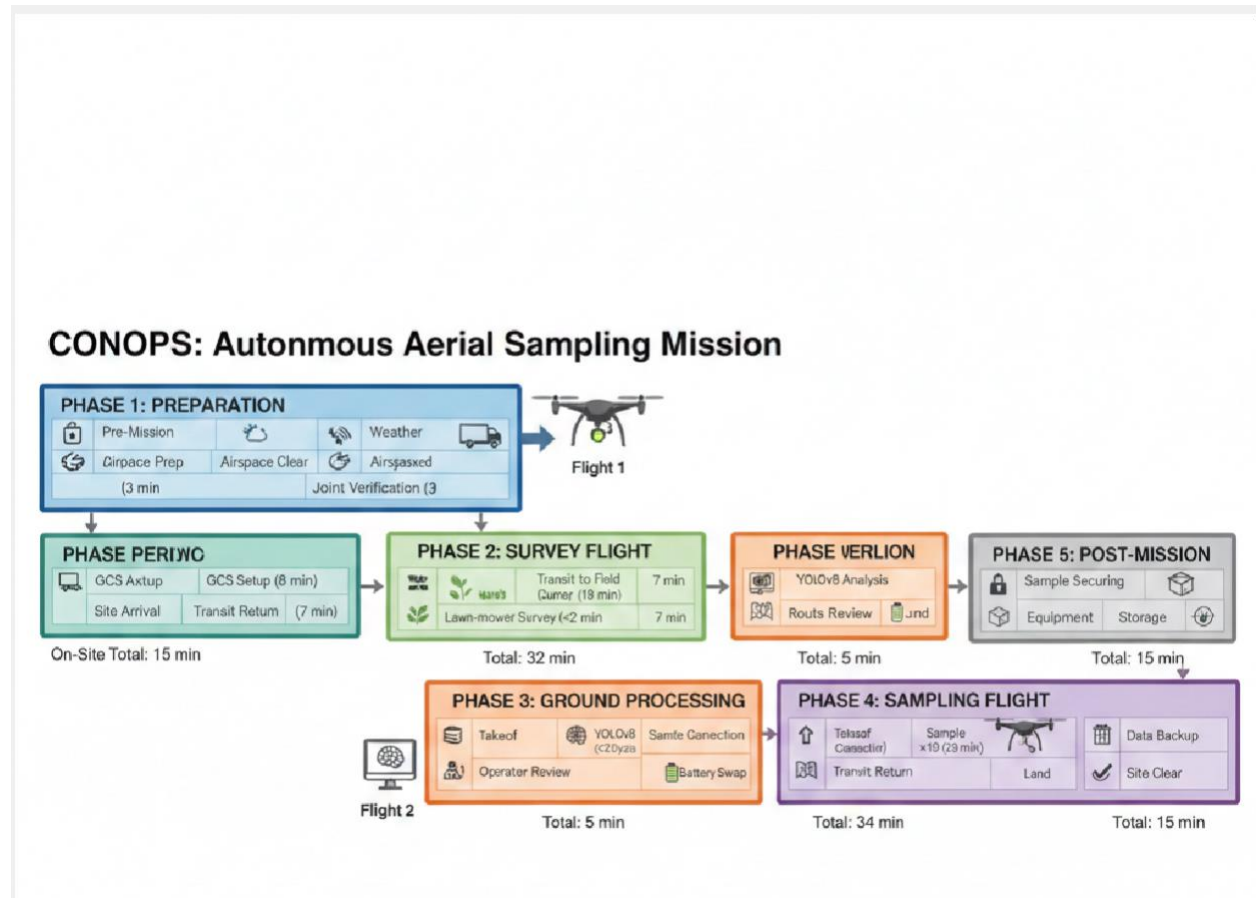


Figure 22. CONOPS mission flow diagram showing complete operational sequence from setup through post-mission

3.1.1 Preparation

Mission preparation begins 24-48 hours before field deployment with administrative tasks: weather forecast review, airspace conflict check (via FAA LAANC if required), landowner coordination, battery charging, and mission planning in desktop software. The pilot reviews the survey pattern for the specific field, confirming waypoint coverage and obstacle clearances. The

payload operator pre-configures AI detection parameters based on the crop type and growth stage.

Ground station setup at the field site takes approximately 8 minutes for the payload operator working alone. The sequence includes: deploy folding tables under canopy tent, position and connect AI workstation and flight control laptop, deploy and orient communication antennas, start generator and verify power to all equipment, launch ground control software, and establish RTK GPS baseline by allowing the base station to collect correction data.

Aircraft preparation proceeds in parallel, taking approximately 7 minutes for the pilot. The sequence includes: remove aircraft from transport case and unfold arms, verify arm lock engagement, install flight battery and connect power, connect detection and sampling payload cables, perform visual inspection of propellers and motors, power on and allow flight controller to complete initialization, verify GPS lock (minimum 10 satellites, HDOP < 2.0), and verify communication links to ground station.

Joint pre-flight verification takes 3 minutes with both operators confirming: telemetry link quality, video feed quality, GPS accuracy (RTK fix or stable float), control surface movement on command, emergency procedures review, and weather conditions remain within limits.

3.1.2 Pest Detection

Flight One executes an automated survey pattern over the target field at 15-20m AGL, traveling at 8 m/s ground speed. The lawn-mower pattern provides 60% lateral overlap between adjacent passes, ensuring complete coverage with no gaps. The RGB camera captures images at 2-second intervals, storing them to onboard SD card for post-flight download.

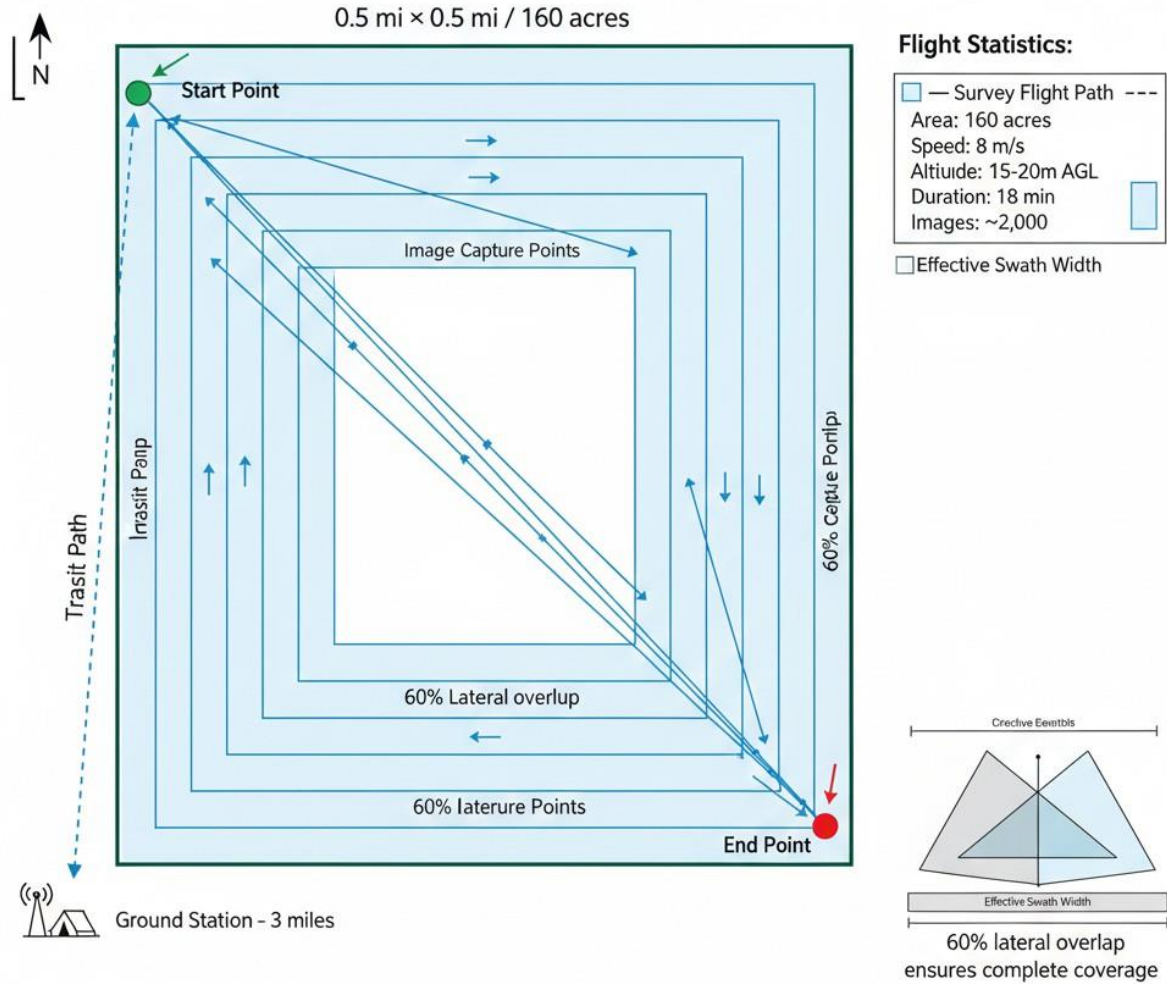


Figure 23. Survey flight pattern showing lawn-mower coverage with overlap zones and image capture points

During Flight One, communication serves situational awareness rather than data transfer. Telemetry (900MHz) transmits position, altitude, battery state, and system health to the pilot's display. Video (5.8GHz) provides real-time imagery to the payload operator for monitoring coverage quality. Voice communication between operators coordinates any real-time adjustments needed.

Flight One lands at the ground station after completing the survey pattern. The pilot removes the SD card from the aircraft while the payload operator begins the AI processing sequence. Image download to the workstation takes 2-3 minutes via USB 3.0 connection. YOLOv8 inference processes the complete dataset in under 2 minutes, generating a map of potential damage locations with confidence scores. The payload operator reviews the top detections, eliminating obvious false positives, and selects up to 10 locations for sampling. The ground station software generates an optimized route using a traveling salesman algorithm to minimize Flight Two duration.

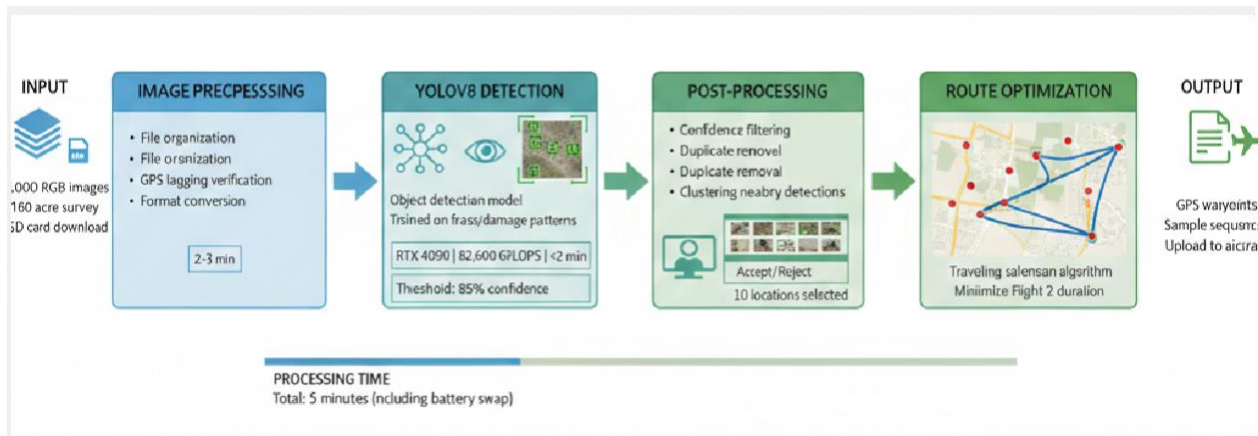


Figure 24. AI processing pipeline diagram showing data flow from image download through route generation

3.1.3 Sample Gathering

Sample location decisions balance AI confidence scores with spatial distribution across the field. The payload operator prioritizes high-confidence detections but ensures sampling covers different field areas to avoid clustering all samples in one zone. The final sample list includes GPS coordinates accurate to centimeter level from the RTK-corrected survey imagery.

Flight Two follows the optimized route to each sample location. At each site, the aircraft approaches at survey altitude, then descends to 3-5m for visual confirmation by the payload operator via HD video. If the operator confirms a valid target, the aircraft descends further to 2-3m sampling altitude, stabilizes in hover, and the operator commands sample collection.

For fruit or leaf samples, the operator selects gripper mode. The positioning gimbal extends the gripper toward the target, the three-finger mechanism closes around the sample, the gripper retracts, and the carousel rotates to present an empty container. For frass samples, the operator selects vacuum mode, positions the collection tube near the frass deposit, activates the pump, and collects material into the inline filter before transfer to a sealed container.

Each sample deposits into a numbered carousel position with automatic GPS and timestamp logging. The system records the carousel position number, GPS coordinates (latitude, longitude, altitude), UTC timestamp, collection mode (gripper or vacuum), and operator notes if any. This metadata creates a complete chain of custody linking each physical sample to its field location.

3.1.4 Post-Mission

After Flight Two landing, sample handling follows a defined protocol to maintain sample integrity for laboratory analysis. The payload operator removes the sample carousel from the aircraft and transfers it to the insulated cooler with ice packs. Each container's seal is verified, and the chain of custody documentation links container numbers to GPS coordinates from the flight log. Samples should reach the laboratory within 24 hours to maintain diagnostic quality.

Data management includes exporting complete flight logs from both flights, backing up all imagery to external storage, and generating a preliminary mission report with coverage maps and detection locations. This report provides growers with immediate actionable information while samples are in transit for laboratory confirmation.

Equipment storage reverses the setup sequence: power down electronics, disconnect batteries (storing in LiPo-safe bags), fold aircraft arms, pack all equipment into transport cases with protective padding, and refuel the generator. A final site inspection ensures no equipment or debris remains at the field location.

3.2 Benchmark Mission

The benchmark mission requires surveying a 0.5 × 0.5 mile (160 acre) field located 3 miles from the ground station and returning 10 samples to the ground station. We analyzed this mission to verify that our system design provides adequate performance margins for successful completion.

Table 14. Benchmark Mission Timeline

Phase	Activity	Duration	Cumulative
Setup	Ground station deployment, aircraft preparation	15 min	0:15
Flight 1 - Transit Out	Climb and transit to field (3 mi @ 12 m/s)	7 min	0:22
Flight 1 - Survey	Lawn-mower pattern (160 acres @ 8 m/s)	18 min	0:40
Flight 1 - Transit Return	Return to ground station	7 min	0:47

Turnaround	Data download, AI processing, battery swap, route planning	5 min	0:52
Flight 2 - Transit Out	Transit to first sample location	7 min	0:59
Flight 2 - Sampling	10 samples @ ~2 min each including transit	20 min	1:19
Flight 2 - Transit Return	Return to ground station	7 min	1:26
Post-Mission	Sample securing, data backup, equipment storage	15 min	1:41
Total Mission Time			1 hr 41 min

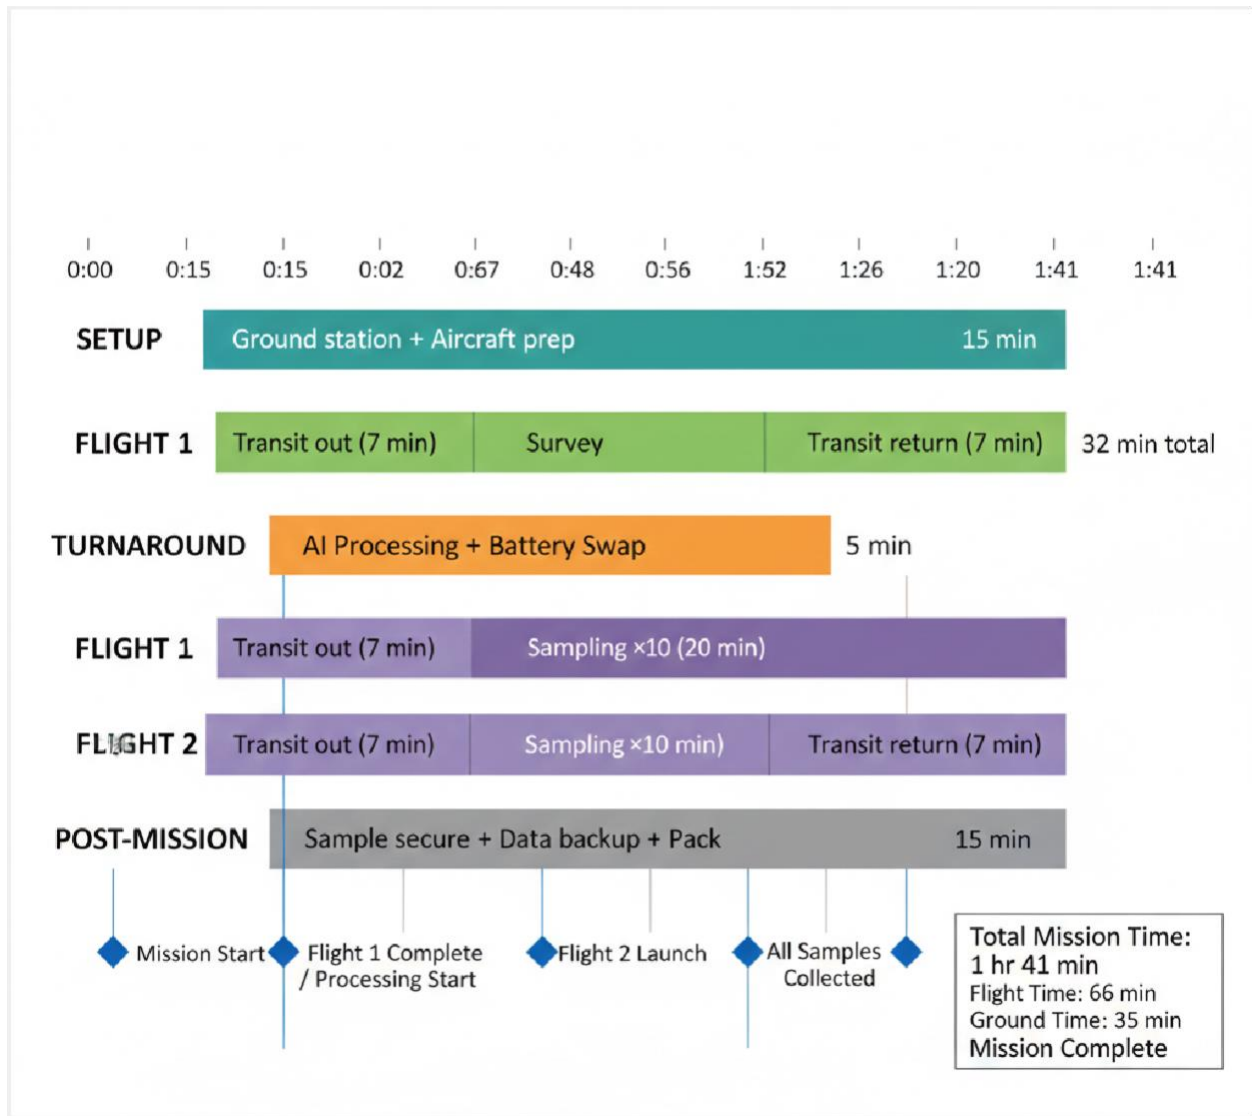


Figure 25. Benchmark mission timeline graphic showing phase durations and critical path

Energy analysis verifies that battery capacity supports each flight with adequate reserve. Our 22,000mAh 6S battery provides 488.4 Wh of stored energy. We calculated energy consumption for each flight phase based on motor power at different flight conditions (hover, cruise, climb) and payload power draw.

Table 15. Flight 1 Energy Budget

Phase	Duration	Avg Power (W)	Energy (Wh)
Takeoff and Climb	1 min	1,200	20.0
Transit Outbound	7 min	720	84.0
Survey Pattern	18 min	650	195.0
Transit Return	7 min	720	84.0
Landing	1 min	600	10.0
Flight 1 Total	34 min		393.0 Wh
Battery Utilization			80.5%

Table 16. Flight 2 Energy Budget

Phase	Duration	Avg Power (W)	Energy (Wh)
Takeoff and Climb	1 min	1,200	20.0
Transit to First Sample	7 min	720	84.0
Sampling Operations	20 min	750	250.0
Transit Return	7 min	720	84.0
Landing	1 min	600	10.0
Flight 2 Total	36 min		448.0 Wh
Battery Utilization			91.7%

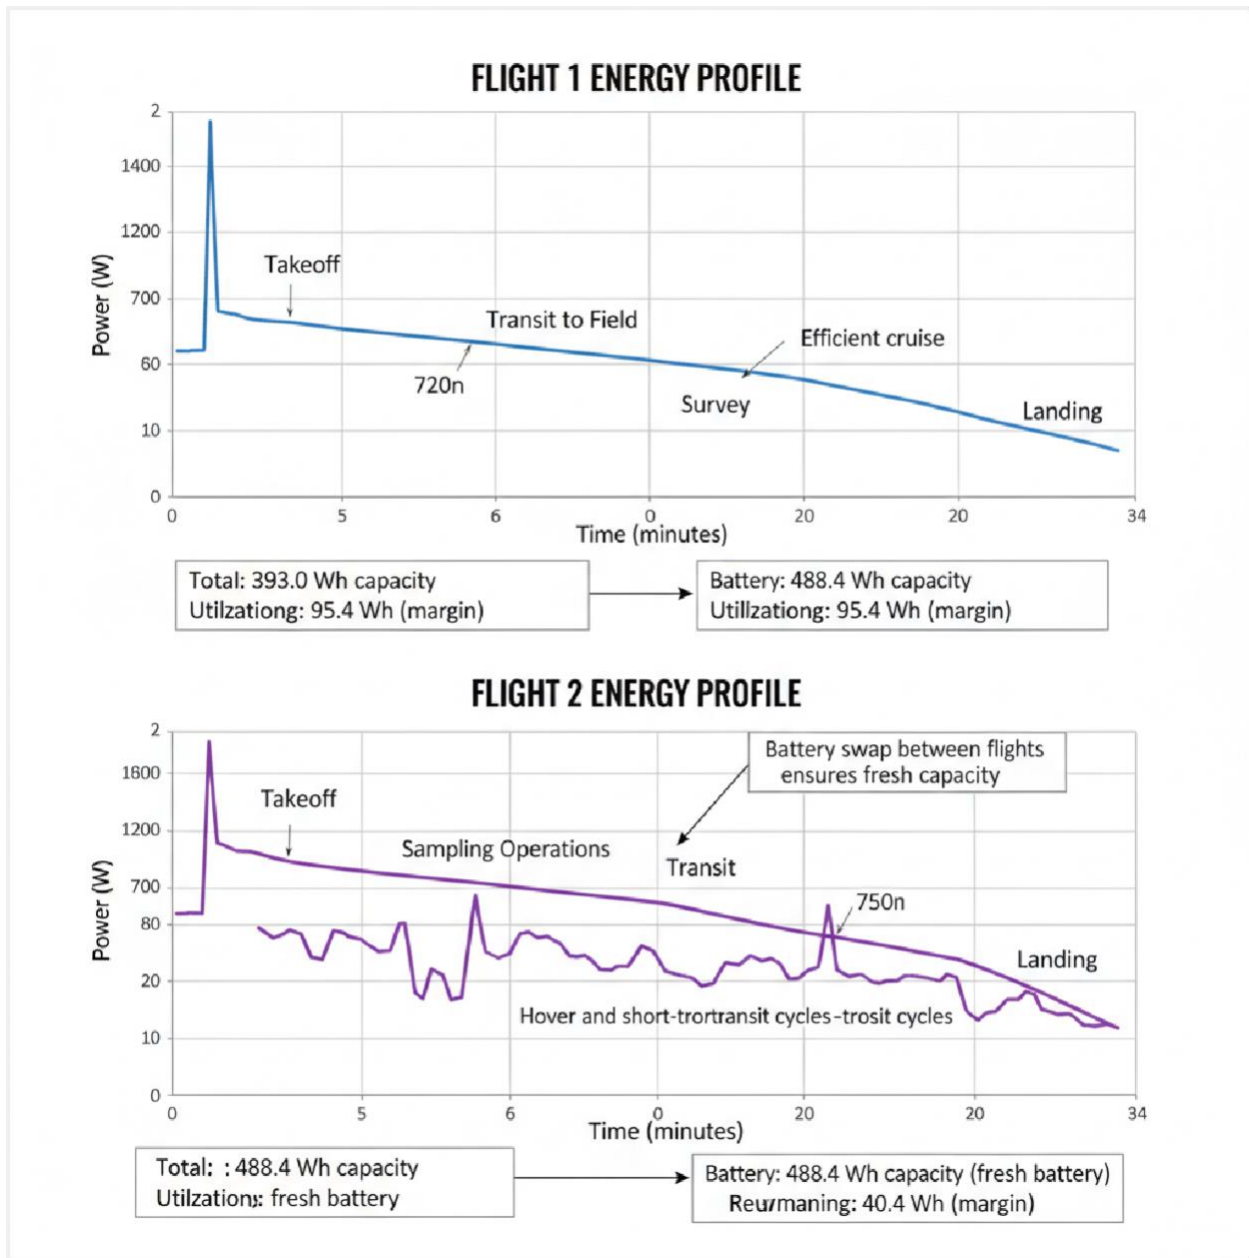


Figure 26. Energy consumption profile showing power draw across mission phases

Flight 1 consumes 393 Wh (80.5% of capacity), providing comfortable margin for weather variations or extended survey requirements. Flight 2 consumes 448 Wh (91.7% of capacity) on a fresh battery. While this margin is tighter, the sampling flight profile includes significant hover time at relatively low power; if battery state drops unexpectedly, the operator can reduce sample count to ensure safe return. The battery swap during the 5-minute turnaround ensures Flight 2 begins with full capacity.

Communication range analysis confirms adequate performance for all benchmark mission phases. Maximum distance occurs during survey operations when the aircraft is at the far edge

of the field (3 miles to ground station plus up to 0.5 miles within the field, totaling approximately 3.5 miles or 5.6 km).

3.3 Safety Requirements

3.3.1 Detect and Avoid

The detect and avoid (DAA) system ensures aircraft and public safety by sensing obstacles, detecting potential conflicts, and executing avoidance maneuvers. Our DAA design addresses both non-cooperative obstacles (trees, structures, birds) and cooperative traffic (transponder-equipped aircraft).

Decision-making for obstacle avoidance occurs primarily onboard the aircraft for time-critical responses, with alerts transmitted to the ground station for pilot awareness. The flight controller processes radar data and initiates avoidance maneuvers automatically if the pilot does not respond within the decision window. This hybrid approach ensures rapid response while maintaining pilot-in-command authority.

Table 17. DAA System Specifications

Parameter	Obstacle Radar (TI AWR1843)	ADS-B In (uAvionix pingRX)
Detection Range	80m maximum	5 nautical miles
Field of View	120° horizontal, 30° vertical	360° omnidirectional
Alert Threshold	30m / 3 seconds to impact	1 nm / advisory
Avoidance Threshold	15m / 1.5 seconds	0.5 nm / warning
Response Action	Automatic climb and hold	Pilot alert, descend if needed

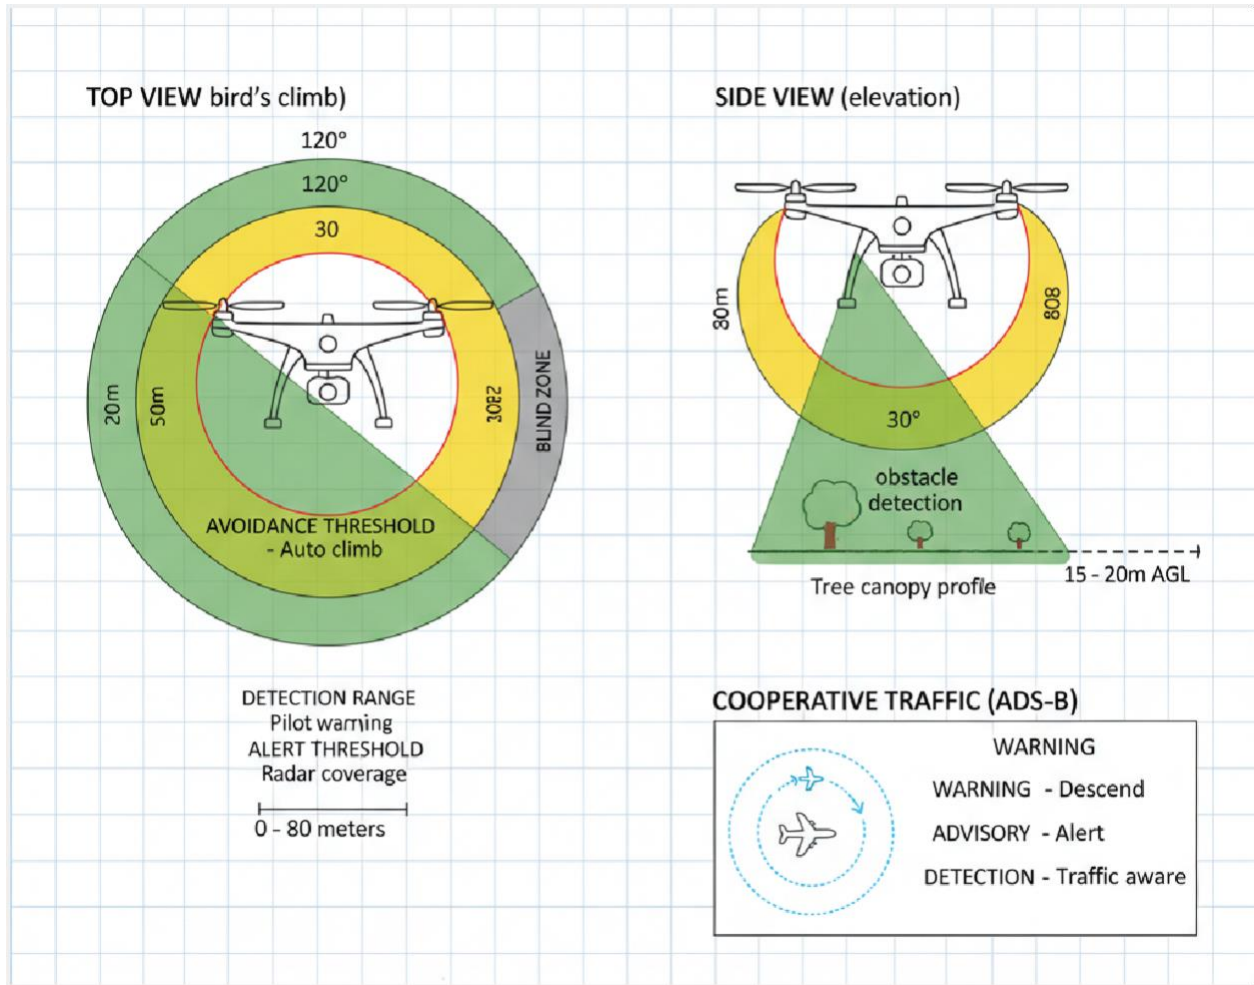


Figure 27. DAA detection envelope showing radar coverage and alert/avoidance thresholds

The 77GHz mmWave radar provides sufficient range and resolution for orchard operations. At survey altitudes of 15-20m AGL, the radar detects tree canopy edges and any obstacles extending above the canopy. At sampling altitudes of 2-3m, the radar detects trunk and branch obstacles. The 30m alert threshold provides approximately 3 seconds of warning at typical approach speeds, sufficient for pilot awareness and automatic system response. The 15m avoidance threshold triggers automatic climb-and-hold if no pilot response is received, ensuring the aircraft does not collide with detected obstacles.

ADS-B In detects cooperative traffic and provides early warning of aircraft entering the operating area. Advisory alerts at 1nm give the pilot several minutes to prepare for potential encounter. Warning alerts at 0.5nm trigger a defensive descent to minimum safe altitude (typically below any manned aircraft approach). These thresholds align with visual line of sight requirements traffic detected at these ranges is also visible to the pilot.

3.3.2 Lost Link Protocol

Lost link protocols define aircraft behavior when communication with the ground station is interrupted. We developed procedures for partial loss (degraded but not absent communication) and total loss (complete communication failure) scenarios.

Table 18. Lost Link Protocol Summary

Condition	Detection Criteria	Aircraft Response	Operator Response
Partial Loss - Telemetry Only	No telemetry 10 seconds	Continue flight, attempt recovery	Verify antenna orientation, check frequency conflicts
Partial Loss - Video Only	No video 5 seconds	Continue if telemetry OK, consider abort	Verify video receiver, switch to telemetry-only if needed
Total Loss	No communication 30 seconds	Climb to 75m, loiter 60 seconds, RTH at 75m AGL	Clear landing area, prepare for automated return
Total Loss - Critical Battery	Total loss AND battery <15%	Immediate landing at current location	Mark GPS position, proceed to recovery

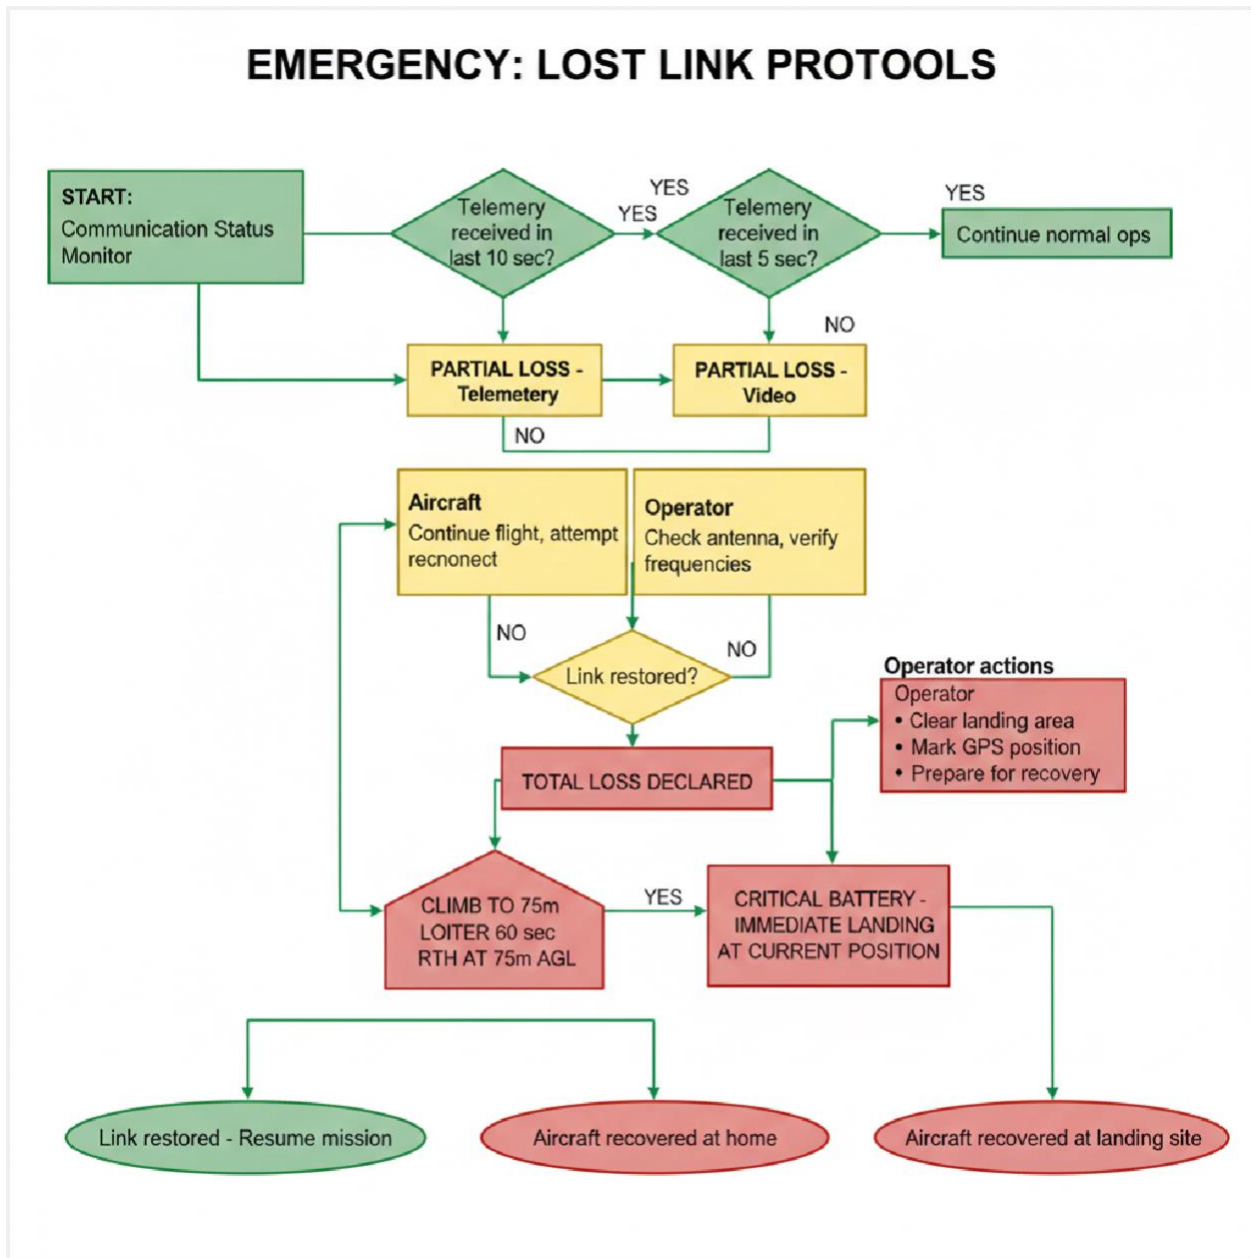


Figure 28. Lost link decision flowchart showing detection, evaluation, and response sequence

The 30-second total loss threshold balances safety against nuisance responses. Brief communication interruptions due to antenna orientation or temporary obstacles should not trigger automated return-to-home procedures. The 75m AGL return altitude provides clearance above most obstacles while remaining well below manned aircraft traffic patterns. The 60-second loiter period allows communication to re-establish before committing to the return flight.

3.3.3 Integration with Manned Aircraft and Other Aircraft

Our UAS operations integrate with manned aviation through compliance with 14 CFR 107 and established right-of-way rules. Per Part 107.37, our UAS must yield right-of-way to all manned aircraft. The pilot maintains visual line of sight to the aircraft throughout operations, enabling direct observation of approaching traffic.

ADS-B In equipment provides electronic traffic awareness to supplement visual observation. Aircraft equipped with ADS-B Out transponders appear on the pilot's display with position, altitude, and velocity. This information enables earlier awareness of traffic than visual observation alone, particularly for aircraft approaching from behind the pilot's position.

If manned aircraft traffic enters the operating area, standard procedure is to descend to minimum safe altitude and hold until traffic clears. Communication with local air traffic control (where applicable) or self-announcing on local CTAF frequencies provides additional deconfliction. Our operations occur in Class G airspace at altitudes below typical manned traffic patterns, reducing encounter probability.

3.3.4 Additional Safety

Additional safety features address emergency scenarios beyond DAA and lost link situations.

Motor Failure: The hexacopter configuration survives single motor failure with degraded performance. If one motor fails, the autopilot automatically reduces power to the opposite motor to maintain balanced thrust, and the aircraft immediately begins return to the ground station. Total thrust capacity drops to approximately 65% of normal, sufficient to maintain controlled flight but requiring immediate landing. A second non-adjacent motor failure permits limited control for emergency landing at the nearest safe location.

Battery Emergency: Battery monitoring provides multiple warning levels. At 20% remaining capacity, the system issues a low voltage warning and the pilot should initiate return. At 15% remaining, the system issues a critical warning and initiates automated return-to-home regardless of current mission phase. Below 10%, the system executes immediate landing at the current location. LiPo-safe storage bags contain batteries during transport and charging.

Weather Deterioration: Operations suspend if sustained wind exceeds 15 mph or gusts exceed 20 mph conditions that compromise hover stability and increase power consumption. Rain or visible precipitation terminates operations immediately due to potential electronic damage and reduced visibility. Lightning detection within 10 miles requires immediate landing and equipment shelter.

Geofencing: The flight controller maintains programmed geofence boundaries defining the authorized operating area. If the aircraft approaches a boundary, it automatically stops and holds position. The 400-foot AGL altitude limit is programmed as an absolute ceiling. Return-to-home will not execute if the planned path would violate geofence boundaries; the aircraft will instead land at its current position.

Personnel Safety: Ground personnel maintain minimum 10m distance from the aircraft during takeoff and landing. No personnel enter the operating area while the aircraft is airborne. The ground station includes a first aid kit rated for lacerations (propeller contact) and burns (LiPo incidents).

4. Business Case

4.1 Cost Analysis

4.1.1 Operating Costs

Operating costs represent the variable expenses incurred to complete each benchmark mission. These costs scale with mission count and include personnel, energy, consumables, and maintenance. We calculated operating costs based on the personnel rates specified in the challenge background document and realistic estimates for other categories.

Table 19. Variable Costs per Benchmark Mission

Category	Calculation Basis	Cost per Mission
Operational Pilot	2.0 hours × \$50.00/hour	\$100.00
Payload Operator	2.0 hours × \$35.00/hour	\$70.00
Electricity (Charging)	1.5 kWh × \$0.20/kWh	\$0.30
Generator Fuel	0.5 gallon × \$4.00/gallon	\$2.00
Battery Cycle Amortization	\$800 battery / 200 cycles × 2 batteries	\$8.00
Sample Containers	12 containers × \$0.50 each	\$6.00
Consumables (Filters, etc.)	Estimated per mission	\$6.00
Maintenance Reserve	Propeller wear, minor repairs	\$20.00
Insurance Allocation	\$1,500 annual / 60 missions	\$25.00
Vehicle Mileage	60 miles × \$0.67/mile	\$40.00
Total Variable Cost per Mission		\$277.30

Personnel costs dominate the variable cost structure at \$170 per mission (61% of total). The pilot and payload operator both work the full mission duration from setup through post-mission, justifying the 2.0-hour allocation for each role. Energy costs are minimal because electric propulsion and generator power are highly efficient. The maintenance reserve of \$20 per mission accumulates funds for propeller replacement (every 50-100 flights) and periodic component servicing.

4.1.2 Fixed Costs

Fixed costs represent the one-time equipment investment required to establish operational capability. We documented costs for all components using current market pricing from established suppliers.

Table 20. Fixed Costs Breakdown

Subsystem	Major Components	Cost
Air Vehicle	Frame, motors, ESCs, flight controller, GPS, batteries	\$2,640
Detection Payload	Sony A7C, MicaSense Altum-PT, gimbal, mounting	\$16,300
Sampling Payload	Gripper, vacuum, carousel, gimbal, controller	\$1,070
C3 Equipment	Radios, video, GPS/RTK, computers, antennas	\$5,980
Ground Support	Cases, generator, charger, tables, accessories	\$3,120
Software Licenses	Mission planning, AI model (open source)	\$180
Total Fixed Cost		\$29,290

The detection payload represents 56% of total fixed cost, driven primarily by the MicaSense Altum-PT multispectral camera at \$13,500. This component provides capabilities beyond minimum requirements teams focused solely on RGB detection could substitute a lower-cost multispectral sensor or eliminate it entirely, reducing system cost by \$8,000-10,000. However, the multispectral capability enhances detection accuracy and provides plant health data beyond pest detection, improving the value proposition for growers.

Each component is necessary to complete the benchmark mission. The air vehicle provides the platform for aerial operations. Detection payload enables damage identification. Sampling payload enables physical specimen collection. C3 equipment enables control, communication, and AI processing. Ground support equipment enables field deployment and sample preservation. Removing any subsystem would prevent mission completion.

4.2 Logistics Details

The benchmark mission requires two personnel throughout the operational period. Each person's role, qualifications, and tasks are defined to ensure mission success without role conflicts or timing impossibilities.

Table 21. Personnel Requirements

Role	Qualifications Required	Primary Tasks	Rate
Operational Pilot	FAA Part 107 certificate, UAS flight experience, safety training	Aircraft setup, flight control, safety monitoring, emergency response	\$50/hour
Payload Operator	AI system training, sensor operation, agricultural knowledge	GCS setup, AI processing, detection review, sample collection commands	\$35/hour

Task timing ensures no conflicts between roles. During setup, the pilot prepares the aircraft while the payload operator deploys the ground station parallel activities that complete within 15 minutes. During flights, the pilot monitors aircraft state and maintains visual contact while the payload operator monitors sensors and AI processing complementary activities with no overlap. During the turnaround between flights, the pilot swaps batteries and conducts post-flight checks while the payload operator manages AI processing and route planning again parallel activities that complete within 5 minutes.

If a single person were required to perform both roles, the mission would take significantly longer and introduce safety concerns. The pilot cannot maintain visual contact with the aircraft while simultaneously monitoring detailed sensor feeds. The payload operator cannot focus on AI analysis while watching for air traffic. The two-person requirement reflects operational necessity, not arbitrary overhead.

4.3 Economic Impact

AGROSCOPE-X creates economic value by enabling precision pest management that reduces both treatment costs and crop losses. We estimated the annual economic impact for a typical 160-acre California apple orchard currently using conventional codling moth management practices.

Table 22. Economic Impact Comparison (160-acre orchard, annual)

Category	Traditional Method	AGROSCOPE-X	Savings
Pest Monitoring	\$5,650 (traps, scouting)	\$3,330 (12 missions)	\$2,320
Pesticide Material	\$43,200 (6 full applications)	\$10,800 (4 targeted, 25% coverage)	\$32,400
Application Labor	\$7,200	\$3,200	\$4,000
Laboratory Analysis	\$0	\$1,800 (sample testing)	-\$1,800
Crop Loss (at \$4.50/lb)	\$28,800 (4% loss)	\$14,400 (2% loss)	\$14,400
System Amortization	\$0	\$5,858 (5-year life)	-\$5,858
Annual Total	\$84,850	\$39,388	\$45,462




Figure 29. Annual cost comparison chart showing traditional vs. AGROSCOPE-X cost breakdown

The analysis assumes AGROSCOPE-X enables conversion from six blanket pesticide applications to four targeted applications covering 25% of the orchard (the portions where damage is actually detected). This 75% reduction in treated area directly reduces pesticide material and application labor costs. Additionally, early detection and targeted treatment reduce crop loss from the industry average of 4% to an estimated 2% the value of knowing where damage is occurring enables intervention before widespread infestation.

The system pays for itself within the first operating season. With annual savings of \$45,462 and a total system cost of \$29,290, the simple payback period is 7.7 months. Even accounting for the time value of money and operational uncertainties, the return on investment is compelling for commercial orchards. Smaller operations could access these benefits through cooperative ownership or contracted services.

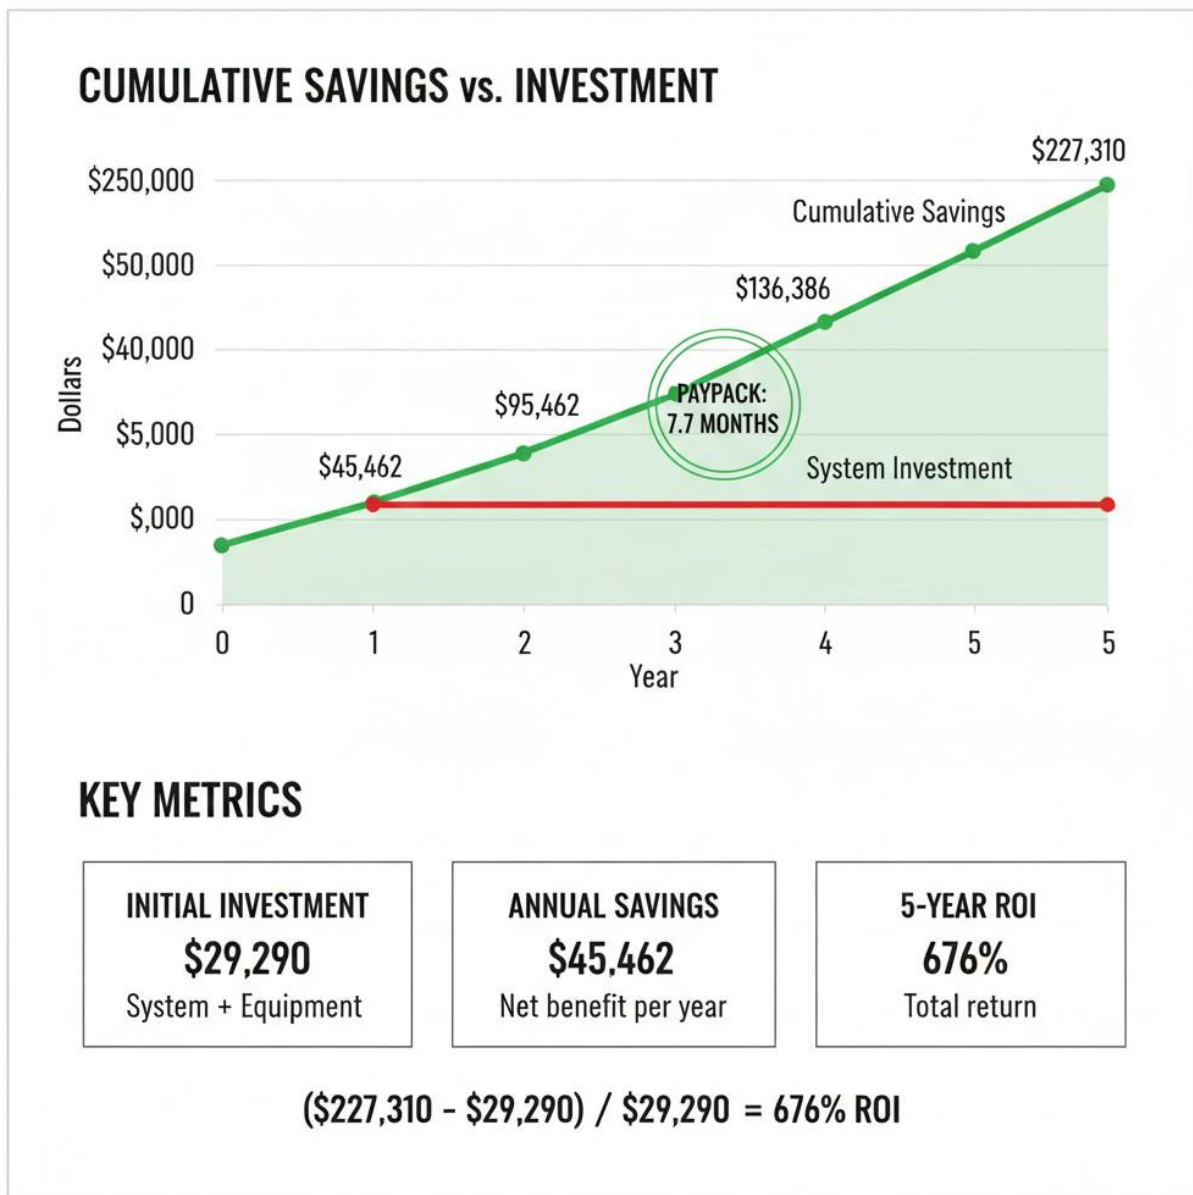


Figure 30. ROI analysis showing payback timeline and cumulative savings over five years

5. Public Affairs/Communications Plan

5.1 Public Relations Strategy

5.1.1 Background and Purpose

AGROSCOPE-X addresses a genuine need in California agriculture, but technical capability alone does not ensure adoption. Our communications strategy addresses the challenges of convincing a traditionally conservative industry to invest in new technology.

Agricultural technology has a mixed history of promises and disappointments. Growers have seen precision agriculture tools that cost more than they saved, required expertise they did not have, or simply failed to work in field conditions. This history creates skepticism toward new solutions, particularly those involving unfamiliar technologies like drones and artificial intelligence.

Additionally, many growers operate on thin margins where capital investments must demonstrate clear return. They need to understand not just that a technology works, but that it will work for their specific operation, with their specific pest pressures, within their specific budget constraints. Abstract claims of efficiency improvement do not resonate as strongly as concrete demonstrations of cost reduction.

Our communications strategy emphasizes demonstrated economic value over technological sophistication. We lead with the business case how the system saves money through reduced pesticide applications and improved crop protection rather than the technical innovation. The two-flight architecture and ground-based AI processing are means to an end (better pest detection) rather than ends in themselves.

5.1.2 Audience and Messaging

We identified three target audiences with distinct information needs and decision-making criteria.

Table 23. Target Audience Messaging Matrix

Audience	Desired Outcome	Key Messages	Channels
Primary: CA Fruit Growers	System adoption or referral	Know WHERE damage is, not just that pests exist. Cut pesticide costs 75%. First-season payback. GPS documentation for compliance.	Trade shows, Extension meetings, demonstration events

Secondary: Agricultural Consultants/PCAs	Recommend to clients	Complements existing IPM. Enables earlier intervention. Provides client reporting data. Reduces callbacks for missed infestations.	Professional conferences, trade publications, direct outreach
Secondary: UC Cooperative Extension	Validate and disseminate	Addresses the detection-to-location gap. Methodology validated by UC researchers. Supports precision agriculture research. Reduces environmental pesticide impact.	Academic presentations, research partnerships

The primary audience California fruit growers makes purchasing decisions based on demonstrated ROI. Our messaging emphasizes the economic benefit ("cut pesticide costs 75%") before the technical capability ("drone-based detection"). Growers care about results: knowing where damage is occurring so they can treat only affected areas rather than entire orchards.

Agricultural consultants and Pest Control Advisors (PCAs) influence grower decisions and specify management programs. They need to understand how AGROSCOPE-X integrates with established Integrated Pest Management (IPM) protocols rather than replacing proven practices. Our messaging emphasizes complementarity the system enhances existing approaches by adding spatial precision to population monitoring.

UC Cooperative Extension provides credibility through independent validation. Extension specialists communicate research findings to growers and can accelerate adoption by endorsing effective technologies. Our connection to Dr. Houston Wilson at UC Riverside provides a foundation for this relationship.


5.2 Products to be Created

Our communications materials include print, digital, and video products tailored to different audiences and distribution channels.

Fact Sheet: A two-page printed document summarizing the problem, solution, and economic case. The front page presents the detection gap problem and our architectural solution in accessible terms. The back page provides system specifications, cost breakdown, and contact information. Designed for distribution at trade shows, Extension meetings, and direct grower contacts.

AGROSCOPE-X
Precision Pest Detection for California Orchards
Team AeroForge | Adrian Wilcox Hill School


AGROSCOPE-X
Team AeroForge



Two-flight system with ground-based AI


THE PROBLEM

- Colling math costs California \$510M annually
- Current traps, slow presence, net location
- Blanket spraying wastes resources




THE SOLUTION

- AGROSCOPE-X
- AI powered detection with 85% accuracy
- Enable targeted treatment



THE SOLUTION

- AGROSCOPE-X maps
- AI powered detection with 95% accuracy
- Enable targeted treatment



SPECIFICATIONS

Aircraft	Hexacopter, 915mm, 5.8kg	
Flight time	32-34 min per flight	
Coverage	160 acres	160 acres
Samples	RGB + Multispectral	
Detection Accuracy	85+	

ECONOMICS

System cost: **\$29,290**


Operating cost: **\$277/mission**

Operating cost: **\$45,462**

Annual savings: **\$45,462**

HOW IT WORKS

- 1 Survey flight captures imagery
- 2 AI identifies damage
- 3 Sampling flight collects specimens
- 4 GPS-tagged samples to lab



KEY BENEFITS

75% less pesticide

50% less crop loss

CONTACT
Team AeroForge
aeroforge.contact@email.com




Figure 31. Fact sheet mockup showing front (problem/solution) and back (specifications/costs) layout

Demonstration Video: A three-minute video showing field operation from setup through sample collection. The video emphasizes operational simplicity ("two people, one hour, ten samples") rather than technological complexity. Includes testimonial content from early adopter growers (to be developed during pilot deployments).

Press Release: A one-page announcement targeting agricultural trade publications (Western Farm Press, Good Fruit Grower, etc.). Leads with the business case, includes system specifications, and provides contact information for follow-up inquiries.

Social Media Content: LinkedIn and Facebook posts with infographics highlighting key statistics ("75% reduction in treated area", "7.7-month payback"). Video clips from demonstration operations. Case study summaries from pilot deployments.




Figure 32. Social media post mockup showing key statistics infographic format

5.3 Distribution Plan

Distribution follows a phased approach aligned with system validation and pilot deployment timelines.

Table 24. Distribution Plan and Costs

Channel	Content	Timing	Lead	Cost
UC Cooperative Extension	Presentati on, fact sheets	Feb-Mar 2026	Ayush	\$0

CA Pear Advisory Board	Meeting presentation	Upon validation	Siddharth	\$0
Trade Publications	Press release	Upon validation	Siddharth	\$0
LinkedIn	Infographics, updates	Biweekly	Aayush	\$200 (promoted)
YouTube	Demo video	Upon completion	Aayush	\$0
Farm Shows	Live demonstration	Growing season	Full team	\$800
Print Production	Fact sheets, handouts	Ongoing	Siddharth	\$300
Video Production	Professional editing	One-time	Aayush	\$500
Total Year 1 Budget				\$1,800

Phase 1 (February-March 2026) focuses on building credibility through UC Extension presentations and industry meetings. These audiences can validate our approach and provide feedback before broader public outreach. Phase 2 (April-May) launches public communications including press release, video release, and social media activation. Phase 3 (June-August) conducts live demonstrations during the growing season when growers can observe actual field performance.

The total Year 1 communications budget of \$1,800 reflects our focus on earned media (press coverage, Extension presentations) and owned media (social platforms) rather than paid advertising. As the system demonstrates validated performance, budget allocation may shift toward expanded farm show presence and additional video content development.

6. Conclusion

AGROSCOPE-X addresses a fundamental gap in agricultural pest management: the difference between knowing that pests are present and knowing where damage is actually occurring. Pheromone traps and degree-day models have indicated codling moth activity for decades, yet growers continue to apply blanket pesticide treatments because they lack spatial damage information. Our system provides that missing information through aerial survey, ground-based AI processing, and precision sample collection.

The two-flight architecture with ground-based processing represents our key technical innovation. This architecture emerged from accepting a fundamental constraint that onboard processing cannot deliver the computational power required for accurate real-time agricultural AI rather than continuing to pursue increasingly compromised single-flight approaches. By separating survey from sampling with ground processing between them, we achieve detection accuracy impossible with onboard systems while maintaining fresh battery capacity for the precision-critical sampling phase.

Every major design decision traces to either mission requirements or expert guidance. The hexacopter configuration provides motor redundancy essential for safe operations among orchard trees. RGB-based detection focuses on frass and entry holes the signatures that Dr. Houston Wilson confirmed as viable for codling moth. Ground-based processing follows Mark Sharp's insight about how successful commercial agricultural operations actually work. The hybrid gripper-vacuum sampling mechanism handles the diverse specimen types that laboratory analysis requires.

The business case demonstrates compelling economics. A \$29,290 system operating at \$277 per mission can save a 160-acre orchard \$45,462 annually through reduced pesticide applications and improved crop protection. The payback period of 7.7 months means growers recover their investment within the first operating season. These economics make precision pest detection accessible to commercial operations rather than limiting it to research demonstrations.

This challenge taught us that constraints drive innovation when approached honestly. Our most significant breakthrough the two-flight architecture came from accepting that we could not make onboard processing work rather than continuing to pursue increasingly marginal solutions. The decisions to eliminate options (not thermal detection, not onboard AI, not single-flight operation) define our design as clearly as the decisions to include components.

AGROSCOPE-X is ready for the benchmark mission and represents Team AeroForge's contribution to precision agriculture in California.

7. References

- California Department of Food and Agriculture. (2024). California agricultural statistics review 2023-2024. Sacramento, CA: CDFA.
- Federal Aviation Administration. (2024). Small unmanned aircraft systems (14 CFR Part 107). Washington, DC: FAA.
- MicaSense. (2024). Altum-PT sensor specifications. Seattle, WA: AgEagle Aerial Systems.
- NASA. (2025). 2025-2026 Dream With Us: High school engineering design challenge background document. Washington, DC: NASA.
- Sharp, M. (2025, December). Personal communication [Video consultation]. Pocatello, ID: Idaho State University.
- Texas Instruments. (2024). AWR1843 single-chip 77GHz mmWave sensor datasheet. Dallas, TX: Texas Instruments.
- University of California Agriculture and Natural Resources. (2024). UC IPM pest management guidelines: Pome fruits. Davis, CA: UC ANR.
- USDA National Agricultural Statistics Service. (2024). Noncitrus fruits and nuts 2023 summary. Washington, DC: USDA.
- Wilson, H. (2026, January). Personal communication [Email correspondence]. Riverside, CA: UC Riverside.