NASA/TM–20260009168

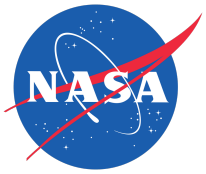

Multi-Modal Site Prediction and Reinforcement-Learned Replanning for Medical UAS Delivery in Montane Forest Regions

* *All authors contributed equally to this research paper. Order listed alphabetically by last name.*

Saadhya Munigela
Thomas Jefferson High School for Science and Technology, Haymarket, Virginia

Rushi Patel
Westford Academy, Westford, Massachusetts

Olivia Przadka
Brooklyn Technical High School, Brooklyn, New York

Anandavel Sakthi
International Academy, Troy, Michigan

Grace Shao
Montgomery Blair High School, Rockville, Maryland

Ian Shelburne
Horace Greeley High School, Chappaqua, New York

October 2026

NASA STI Program Report Series

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at `http://www.sti.nasa.gov`.

NASA/TM–20260009168

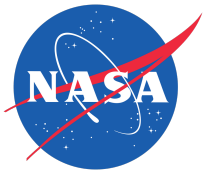

Multi-Modal Site Prediction and Reinforcement-Learned Replanning for Medical UAS Delivery in Montane Forest Regions

* *All authors contributed equally to this research paper. Order listed alphabetically by last name.*

Saadhya Munigela
Thomas Jefferson High School for Science and Technology, Haymarket, Virginia

Rushi Patel
Westford Academy, Westford, Massachusetts

Olivia Przadka
Brooklyn Technical High School, Brooklyn, New York

Anandavel Sakthi
International Academy, Troy, Michigan

Grace Shao
Montgomery Blair High School, Rockville, Maryland

Ian Shelburne
Horace Greeley High School, Chappaqua, New York

National Aeronautics and
Space Administration

NASA Ames Research Center
Moffett Field, California 94035

October 2026

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA STI Program / Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

Acknowledgments

This research would not have been possible without the guidance and expertise of our mentors, Estela Buchmann and Priyank Pradeep. We also thank the NASA Ames Aviation Systems Division for the opportunity to pursue this work.

Contents

1 Nomenclature **3**

2 Introduction **4**

3 Background and Motivation **4**

4 Problem Formulation **5**

5 Assumptions **5**

6 Vehicle Dynamics and Energy Framework **7**
- 6.1 Vehicle Selection and Configuration . . . 7
- 6.2 Power Model . . . 8
- 6.3 Mission Energy . . . 11

7 Pre-Flight: Geospatial Site Prediction **11**
- 7.1 Framework Overview . . . 11
- 7.2 Study Regions and Data . . . 12
- 7.3 Label Generation: The Lidar Oracle . . . 13
- 7.4 Model Architecture and Training . . . 14
- 7.5 Site Ranking and Energy-to-Service . . . 15

8 In-Flight: Deep Reinforcement Learning Agent **16**
- 8.1 Framework Overview . . . 16
- 8.2 The Reinforcement Learning Framework . . . 17
  - 8.2.1 Simulation Environment . . . 17
  - 8.2.2 What the Agent Sees and What It Can Do . . . 18
  - 8.2.3 Reward Function . . . 18
  - 8.2.4 Learning Algorithm . . . 19
  - 8.2.5 Onboard Perception and Sensor Fusion . . . 19

9 Results **20**
- 9.1 Vehicle and Energy Results . . . 20
- 9.2 Pre-Flight Model Performance . . . 21
- 9.3 Cross-Biome Transfer and Label Efficiency . . . 22
- 9.4 In-Flight Policy and Mission Performance . . . 24

10 Conclusions **26**

11 Future Work **26**

Appendix **28**

References **29**

1 Nomenclature

Table 1: Symbols and abbreviations used in this report

Symbol	Description
Mathematical Symbols	
$\mathbb{1}[\cdot]$	Indicator function, equal to one when its predicate is true
r_{winch}	Required clear-column radius for winch delivery
L_{cable}	Maximum available winch-cable length
E_{service}	Modeled energy required to service a candidate site
m	Total aircraft mass
g	Gravitational acceleration
ρ	Air density
f_{par}	Equivalent flat-plate parasite area
C_{d0}	Blade profile drag coefficient
σ	Rotor solidity
V_{tip}	Rotor tip speed
R	Rotor radius
A_{ind}	Effective rotor-disk area for induced power, four coaxial discs
A_{pro}	Blade-swept area for profile power, eight rotors
k_{μ}	Advance-ratio factor in the profile-power expression
P_{avionics}	Continuous avionics and sensor hotel load
P_{hover}	Electrical hover power at the analysis mass
P_{compute}	Onboard inference load charged to the scan phase
e	Electrical energy per unit distance at a given airspeed
d	One-way ground distance from launch origin to a candidate
t_{scan}	Duration of the descending scan at a candidate site
t_{winch}	Duration of one winch out-and-back cycle
h_{surr}	Neighborhood-maximum canopy height around a site
h_{clr}	Vertical clearance held between aircraft and canopy
w_{gap}	Measured canopy-opening width
κ	Induced-power correction factor
η_{prop}	Propulsive-system efficiency
v_{winch}	Winch line speed
r_t	Intermediate reward at time step t
E_{usable}	Usable battery budget after the reserve is held back
ΔE_t	Energy spent by the action taken at time step t
R_{success}	Reward for a completed delivery
w_{energy}	Energy penalty weight, derived from R_{success}
C_{crash}	Penalty for true battery exhaustion
C_{reserve}	Penalty for finishing below the energy reserve
C_{abort}	Penalty for aborting a mission
$\mathbb{E}[\cdot]$	Expectation operator
E_{total}	Total mission energy
p_{viable}	Calibrated probability that a candidate is viable
ϵ_p	Probability floor used in expected-energy ranking
r_{land}	Required clear-disc radius for landing
A_{rotor}	Disk area of one effective rotor
Ω	Rotor angular speed
E_{batt}	Installed battery energy
V_{max}	Airspeed envelope limit used for the power curve
Abbreviations & Acronyms	
AUPRC	Area under the precision–recall curve
ECE	Expected calibration error
LiDAR	Light detection and ranging
RGB	Red–green–blue optical imagery
RL	Reinforcement learning
UAS	Uncrewed aircraft system
VTOL	Vertical takeoff and landing

2 Introduction

In 2016, an estimated 3.6 million people in low- and middle-income countries died because they did not get health care they needed [1]. Part of that gap comes down to transport. In mountains and forests, roads are slow, broken, or missing, and medical supplies do not arrive.

Delivery drones can fly over the terrain that roads cannot cross. But today's drone networks only fly to places that have already been surveyed on the ground and recorded as approved delivery sites. If a community has no such site, the network cannot serve it, no matter how badly supplies are needed. The drone can make the flight. What it lacks is a known place to put the package.

This report describes a system that finds that place. It has two parts. Before launch, a model reads free satellite and elevation data over a region and produces a ranked list of candidate delivery sites. Each candidate carries a viability score (the probability that a delivery there would succeed), a delivery mode (land or winch), an estimated energy cost, and a confidence flag. The flag marks candidates the model is unsure about, so the aircraft knows to verify them from the air before trusting them.

The aircraft then flies to the top-ranked site. On arrival, an onboard perception module combines visible, thermal-infrared, and lidar depth data and decides in real time whether the site is clear. If the site is blocked, a reinforcement-learning agent picks the next candidate and re-plans against the remaining battery. Flying home is the last resort. The agent can afford to divert because it already knows where the other candidates are and what each would cost to reach.

3 Background and Motivation

About 56% of the world's rural population cannot get essential health services [2]. Cost, staff shortages, and supply shortages all play a part, but so does the difficulty of moving supplies to where they are needed. In a review of 43 studies of emergency obstetric care in developing countries, drug procurement and logistics were named as a barrier in 65% of them, second only to inadequate training at 86% [3]. Poor rural roads are a documented cause of gaps in medicine distribution [4].

The supply problem is worst where roads are hardest to build and keep. About 915 million people, roughly 13% of the world, live in mountains [5, 6], and about 1.6 billion rural people lived within 5 km of a forest in 2012 [7]. In mountains, steep slopes make road travel slow or impossible. In forests, fallen trees, ravines, and undergrowth block ground routes, and rivers only go where the water goes.

Cargo drones already fill part of this gap. Zipline, the largest medical drone operator, passed two million deliveries by January 2026 [8], across seven countries and several U.S. markets, including a partnership with the Cleveland Clinic [9]. Its first aircraft parachutes a package into a cleared area about the size of two parking spaces. Its second lowers a sensor-equipped "droid" on a tether that checks the ground is clear and settles onto a spot as small as a meter across [10, 11]. Checking a site on arrival is therefore a capability that already works in the field.

The limit of these systems is that the destination is chosen before takeoff. If the site turns out to be blocked, the aircraft aborts and flies home [12]. For repeat deliveries to fixed, surveyed clinics this is the right design. But where no surveyed destination exists, the aircraft cannot be sent at all.

How the package leaves the aircraft matters too. An airdrop needs an open area, and tree cover or wind drift defeats it. A landing needs a flat patch several meters across, and slope rules it out quickly. A winch needs only a narrow vertical gap through the canopy, and because the aircraft never touches down, slope matters far less. Ground that is both steep and covered by closed canopy has no conventional delivery method at all. Mountain forests are that combination.

Three facts motivate the work. Aircraft save a lot of time in this terrain: in the mountains of northeast India a 50 km road trip of three to four hours was replaced by an 18 km flight of 20 minutes [13]; drone delivery in the Hainan mountains roughly halved delivery times [14]; and a delivery to a remote Amazon village 40 km away took 35 minutes by air against about six hours by riverboat [15]. That time saves lives at the supply end. A ten-year review of maternal deaths at a regional hospital in Egypt found that delay

inside the facility, driven largely by shortages of blood and essential drugs, played a role in 88.9% of cases [16]. Those patients had already reached care and were waiting on stock, which is what a small, frequent aircraft delivers well. And the cost is low. A single helicopter medical transport was billed at a median of about $36,400 in 2017 [17], which is the charge for one flight. The aircraft we model has a list price of about $12,500 [18], a one-time purchase spread across every flight it makes. Per delivery, one comparison put drone delivery in remote areas at about $0.88 against $2.22 by ground [19].

None of this reaches a location without a registered delivery site. Registering one means a ground visit to terrain that is hard to reach on the ground, and the record goes stale as vegetation grows back. Having the aircraft search for a site on arrival means surveying while hovering, which is the most expensive way it can spend its battery. What is missing is knowledge, before takeoff, of where a package could plausibly be placed.

4 Problem Formulation

Existing delivery systems assume the destination is known, surveyed, and registered before takeoff. That assumption breaks in several common situations. The person who needs the delivery may not be at a facility at all, for example a landslide casualty or a stranded group. A disaster may make registered facilities unsafe or unreachable just when demand peaks. And across much of the terrain described above, there is simply no registered facility within range. In each case the aircraft could make the flight. The mission fails because there is no valid destination to plan against.

The problem we address is choosing the delivery site. It is chosen once before launch and checked again during the descent, and at both points the question is the same: where in unprepared terrain can a package be placed?

We split the work into two stages because of what each kind of data costs. Satellite and elevation data cover the whole planet and are free, but they are coarse. They cannot resolve a clearing a few meters wide, and the images may be months old, so nothing computed from them can be treated as a final landing point. Sensors on the aircraft can resolve the clearing, but only from close range and only while the aircraft hovers and burns battery. So the pre-flight stage proposes sites and the in-flight stage verifies them.

The pre-flight stage output is a ranked candidate list passed as input to the in-flight stage. Each candidate entry is a map coordinate with the modeled probability that a delivery there succeeds, whether it would be a landing or a winch drop, the energy it would cost to serve, and whether the model was confident enough to be trusted without a look. In flight, the aircraft goes to the highest-ranked candidate and checks it during descent. If it is clear, the aircraft delivers. If not, it moves to the next candidate, whose position, likelihood, and cost are already known. The aircraft picks among a few known options instead of searching open ground from a hover.

This report selects destinations and nothing else. Four neighboring problems are treated as assumptions: planning the flight path, keeping that path clear of obstacles, keeping the path and descent column clear of wires and power lines, and knowing who needs a delivery and when. They are the first four rows of Table 2. One more restriction is a design choice: the pre-flight stage never uses live sensing, which is what lets it run for any region ahead of time.

5 Assumptions

Table 2 lists every assumption the rest of the report depends on. The first four fix the scope. The rest are modeling simplifications, each justified where it is used.

Two entries have effects beyond the table. Hover power grows as air gets thinner: at 2,000 m, holding the same mass in a hover takes about 10.3% more electrical power, and that term enters the energy formula of Section 6 directly. Also, the regulatory assumption depends on rules that are still changing. These missions fly beyond the pilot's line of sight by their nature. Under current Part 107 rules [22] that requires a case-by-case waiver. The proposed Part 108 framework would replace those waivers with standardized, risk-based approval [23]. It was directed by Executive Order 14307 [24], required by the FAA Reauthorization Act of

2024 [25], and entered interagency review in July 2026. It is not yet final.

The modeling assumptions all lean the same way. Sea-level air density, a constant induced-power factor, a calm weather window, and a clear atmosphere all make the energy budget and the reachable area look better than real flight would, and all but the second get worse as elevation and canopy increase. So the model is most optimistic in the terrain this report is about.

Table 2: Assumptions carried by the analysis

Area	Assumption	Effect on results
Trajectory (scope)	The route from the launch point to a chosen site is known. We do not plan it here.	Site selection is judged separately from the route flown
Obstacles (scope)	The path to a chosen site has no obstacles.	No en-route avoidance is modeled
Wires and power lines (scope)	No wires or power lines lie along the path or in the descent column.	Detecting them, which any real descent into unprepared ground would need, is out of scope
Demand (scope)	Who needs a delivery, and when, is given. Section 11 discusses modeling it.	The report answers where a package can be placed, and does not ask where one is needed
Atmosphere	Air density is fixed at the sea-level value of 1.225 kg/m^3 everywhere.	Understates energy at altitude by about 10% at 2,000 m
Weather	Flight is approved or refused before dispatch. Wind enters as a constant from regional climate data.	Removes gusts and headwind penalties
Flight mechanics	Each flight phase is steady. The aircraft is a point mass at fixed heading and constant airspeed.	No rotor flapping, control lag, or gust response
Rotor inflow	The induced-power correction factor is 1.28 in hover, cruise, and transition alike.	Weakest in transition; understates induced power
Payload and reserve	The payload is 5.0 kg. One fifth of the battery is held back unused.	Sets the analysis mass and the derated range
Bare-earth terrain	The FABDEM elevation product is treated as exact.	Its 2.88 m error in forest, against 5.15 m for the Copernicus GLO-30 model it corrects [20], is not carried into slope or clearance
Lidar ground model	Ground returns under the canopy are dense enough to build a reliable bare-earth surface.	Published errors of 0.32–1.22 m [21] are not carried forward
Satellite input	The imagery is a leaf-on seasonal composite.	A composite of a typical season, never a live scene
Label validity	The lidar-derived labels stand in for physical viability.	No site they produced was ever flown
Onboard sensing	The air is clear enough for all three sensing channels.	All three work in optical and infrared bands, so they degrade together

Continued on next page

Table 2 – continued from previous page

Area	Assumption	Effect on results
Verification	Onboard sensing catches the pre-flight stage's false positives.	Justifies the recall-first training loss
Regulation	Flight beyond visual line of sight is legal under the proposed Part 108 framework.	Currently requires case-by-case waivers

6 Vehicle Dynamics and Energy Framework

A site is only "viable" relative to a specific aircraft that has to reach it, hover over it, and lower a package into it. So this section fixes one reference aircraft and builds the energy model the site ranking uses. Three of its outputs are handed forward: the clear radius the aircraft needs beneath it to run the winch, the length of cable it carries, and the electrical energy it spends per meter of cruise. Changing the aircraft would change the learning problem, which is why the label definition in Section 7 was kept free of any vehicle parameter.

6.1 Vehicle Selection and Configuration

An energy model needs published mass, rotor geometry, battery capacity, and endurance. The operators with the largest delivery fleets publish none of these, because they build and fly their aircraft in house. Amazon is the only partial exception, and its FAA filing gives mass but no aerodynamic or sensor data [26]. So we limited the candidates to commercially available aircraft with a complete published specification. Fourteen were reviewed against six criteria: a complete specification; sensing that can resolve terrain; an integrated winch; the ability to sense and deliver on the same flight; enough payload margin left over after both are fitted; and a supported way to run computation onboard. Two criteria eliminated most of the field. Several aircraft offer both a lidar option and a payload mechanism, but the two share one payload allocation, so the aircraft has to choose between seeing a site and serving it. And the small survey platforms carry under 1 kg before a winch is fitted, far below our 5 kg reference payload.

Table 3 lists the aircraft reviewed, from manufacturer specification sheets [27]. The payload column is capacity before winch and sensing hardware are subtracted, so every figure overstates what a package could weigh. We selected the DJI FlyCart 100, listed first. It was the only aircraft reviewed that publishes a complete specification and carries both factory-fitted lidar and an integrated winch installed without exceeding the payload margin.

Table 3: Candidate aircraft reviewed

Aircraft	Type	Payload	Terrain sensing	Reason not selected
DJI FlyCart 100	Coaxial quad	85 kg	Lidar, mmWave radar	Selected; not NDAA-compliant, closed flight software
Wingcopter 198	Tilt-rotor VTOL	4.7 kg	Ground altimeter	Lidar and payload mutually exclusive; no winch
Draganfly Heavy Lift	Multirotor	30 kg	Ground altimeter	Lidar and payload mutually exclusive
Harris Aerial Carrier H6	Hexacopter	8 kg	Lidar	Lidar, thermal, and winch mutually exclusive
Vision Aerial Vector	Hexacopter	5 kg	Lidar	Requires a separate compute module

Continued on next page

Table 3 – continued from previous page

Aircraft	Type	Payload	Terrain sensing	Reason not selected
Quantum Systems Trinity F90+	Fixed-wing hybrid	0.7 kg	Lidar	Payload below delivery mass
DJI Matrice 400	Multirotor	—	Lidar by add-on	No integrated delivery mechanism
Skydio X10	Multirotor	0.5 kg	None	Payload below delivery mass; no lidar
ModalAI Starling 2 Max	Multirotor	0.5 kg	None	Payload below delivery mass; drop mechanism undesigned
Freefly Astro	Multirotor	—	Lidar, base station required	Lidar and cargo options not offered on one airframe
A2Z Air Dock	Multirotor	—	Sonar, RTK, vision	No lidar; sonar useful only near the ground; separate compute required

That selection comes with caveats about availability. The FlyCart 100 does not comply with the American Security Drone Act [28], and since 23 December 2025 its manufacturer has been on the FCC Covered List, which blocks new models from the authorization needed for U.S. import and sale [29, 30]. Its flight software is also closed, which matters because the decision agent of Section 8 would have to run inside it. The FlyCart 100 is therefore a modeling reference chosen for the completeness of its specification. It is not a test article.

The FlyCart 100 is a coaxial four-axis multirotor: four arms, each carrying a pair of 62 in (1.5748 m) propellers spinning in opposite directions one above the other, for eight rotors on four axes. Table 18 in the Appendix collects every configuration value the energy model uses, with its source. At the analysis condition the aircraft weighs 94.6 kg: 60.2 kg of airframe and winch, two 14.7 kg batteries, and a 5 kg payload. The winch footprint radius of 3.0 m is the one entry that passes directly to the pre-flight stage. It is the clear radius needed beneath the aircraft: the payload's own radius plus an allowance for the payload swinging on the cable plus an allowance for imperfect position hold. The 2.10 m swing allowance is an assumption, and because this radius defines every training label, the whole label set rests on that assumption.

The airframe ships with lidar, five optical cameras, a GNSS receiver, and a payload weight sensor [18]. The lidar is what suits it to forest work, because unlike a camera it returns from thin obstacles such as branches. We add two payloads. An NVIDIA Jetson Orin Nano Super [31] runs the verification model onboard, since no link to a ground station can be assumed over this terrain; its 25 W is charged to the scan phase only. And because the stock cameras report relative brightness rather than calibrated temperature, a Teledyne FLIR Hadron 640R+ supplies both the visible and thermal branches of the perception module from one 56 g, 1.8 W unit [32], through the manufacturer's payload interface [33].

6.2 Power Model

Every energy figure in this report follows from how much power the aircraft needs at a given airspeed. We model the aircraft as a point mass in steady, level, unaccelerated flight at constant airspeed. Gross weight is held constant, there is no climb, and wind is ignored. The rotors tilt forward so that the vertical part of their thrust carries the weight (928.0 N at the analysis mass) and the horizontal part balances drag. At the 15.8 m/s cruise speed that tilt reaches 13.3°, which is large enough that we solve the balance exactly instead of using a small-angle approximation.

The shaft power the rotors must deliver is the sum of three terms, and we keep them separate because they move in opposite directions as speed changes. Induced power is the cost of pushing air down to make lift. It dominates in hover and falls in forward flight, because the rotors keep meeting fresh air instead of recirculating their own downwash. We compute it from Glauert's forward-flight inflow relation [34] with a correction factor of 1.28 for non-uniform inflow, tip losses, and the interference between the upper and lower rotor of each coaxial pair [45]. Parasite power is the cost of dragging the airframe through the air. It grows with the cube of airspeed, so it is negligible in hover and dominant at high speed, and it is the reason the

aircraft has a most-efficient speed at all. Profile power is the drag on the blades themselves. It is nearly constant because rotor speed is fixed, and it is the smallest of the three at cruise.

Three properties of this airframe are not published and could not be measured, so we derived each from an independent source. Fitting them to the aircraft's published range would have made any later range comparison circular. The airframe's equivalent flat-plate drag area is 1.44 m^2, from the published body dimensions and a bluff-body drag coefficient at the top of the range established for cubes, flat plates, and open multirotor frames [35, 36, 37, 38, 39]. This value is probably low, because it ignores the extra frontal area the aircraft presents when tilted and the drag added by rotor downwash [40, 41, 42]. The blade drag term comes from T-MOTOR propeller bench data [43, 34, 44]. The induced-power factor comes from the coaxial interference model already cited.

Shaft power becomes a battery draw by dividing by a propulsion efficiency of 0.85 and adding a 78.5 W avionics load. We swept airspeed from hover to 23 m/s in MATLAB, solving the inflow relation at each point. Two speeds come off that curve. The maximum-endurance speed, 11.1 m/s, minimizes power and so maximizes time aloft. The maximum-range speed, 15.8 m/s, minimizes energy per meter and so maximizes distance on one charge. We use the maximum-range speed everywhere below.

Table 4 gives the shaft power breakdown at cruise. Induced and parasite power are within two percentage points of each other, which puts the cruise speed close to where the two halves of the curve cross. Dividing the 8021.0 W shaft total by the efficiency and adding the avionics load gives 9514.9 W electrical, and dividing by 15.8 m/s gives 602.2 J/m. So every kilometer of straight-line travel distance to a candidate delivery location costs 602.2 kJ each way.

Table 4: Shaft power split at the cruise condition

Component	Mechanical power (W)	Share of total (%)
Induced power, the cost of making lift	3643.1	45.4
Parasite power, the cost of airframe drag	3478.9	43.4
Profile power, the cost of blade drag	898.9	11.2
Shaft power required	8021.0	100.0

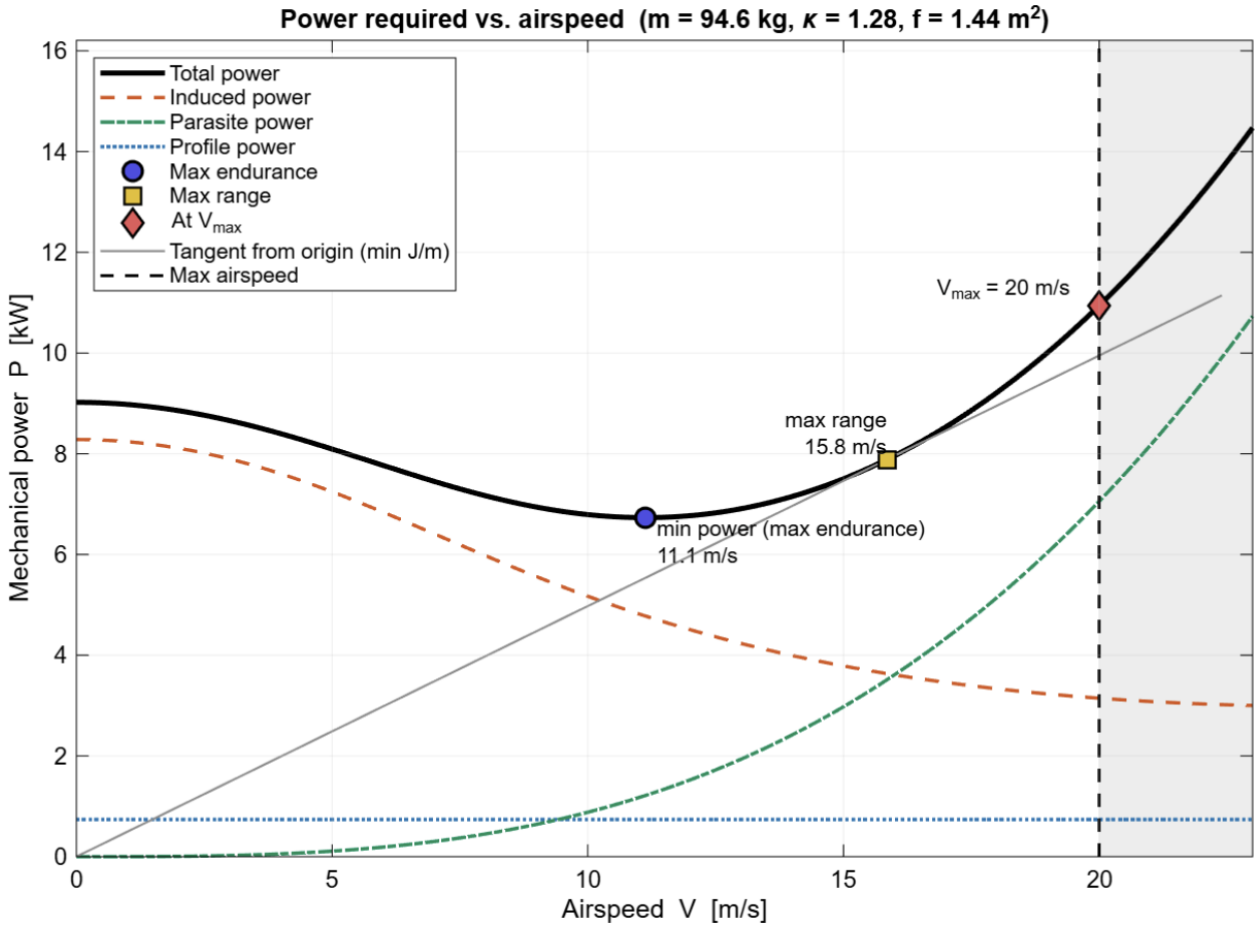

Figure 1: Power required against airspeed

Every point on Figure 1 is mechanical shaft power at the 94.6 kg analysis mass. The profile term is plotted at its hover value, which understates the plotted total at cruise by about 2%; the tabulated value includes the correction. The curve is shallow between the marked speeds for max endurance and max range, indicating that power is less sensitive to exact cruise speed in this range.

None of the three derived parameters was fitted to published performance, so integrating the power curve is an independent prediction of range. The predicted range is 23.9 km at the analysis mass against about 26.0 km published at empty weight, a reasonable gap given the mass difference.

Three limitations remain. The first is that the cruise speed depends on which version of the model produces it. Three sources give three different answers. An earlier version of this curve, computed before the three derived parameters above were revised, put the maximum-range speed at 16.9 m/s. The current version puts it at 15.8 m/s. And the published 12 km in 14 minutes implies an average of 14.3 m/s. Each speed carries the energy intensity computed with it, so a speed from one version must never be combined with an energy intensity from another. The second is that the inflow correction is held constant across flight regimes it was not calibrated for. The third is that none of this has been verified against live flight data from this vehicle.

6.3 Mission Energy

Reaching a site is only part of its cost. The aircraft then has to hold position and run the cable out and back. We chose not to include a separate winch-motor energy cost because the manufacturer did not publish one and because the estimated power required to raise a 5 kg payload for 50 s at 1.2 m/s is negligible compared to the power required for the aircraft to hover during that time. The hover power itself, 10.69 kW, is derived from the manufacturer's published hover endurance scaled to the 94.6 kg analysis mass. We did not derive the pendulum dynamics that set the footprint, so $r_{\text{winch}} = 3.0$ m rests on the assumed swing allowance, and that assumption flows into every delivery site count this report gives.

There are two ways to run a winch delivery, and we use both, because a site that fails one may pass the other. In the first, the aircraft hovers above the canopy and lowers the package through the opening, so the cable has to reach past the tallest tree around the opening, not the average one. In the second, the aircraft descends into the opening first, so the constraint is lateral clearance:

$$L_{\text{cable}} \;\geq\; h_{\text{surr}} + h_{\text{clr}} \quad \text{(hover above canopy)}, \qquad w_{\text{gap}} \;\geq\; 2\, r_{\text{winch}} \quad \text{(descend into opening)}. \tag{1}$$

Here $L_{\text{cable}} = 30$ m is the installed cable length, h_{surr} is the maximum canopy height in the neighborhood of the site, h_{clr} is the clearance the aircraft keeps above the canopy, and w_{gap} is the measured width of the opening. Both conditions are applied as filters at ranking time rather than in the labels.

The energy to serve one site has a cruise part, a station-keeping part, and a compute part:

$$E_{\text{service}}(d) = 2\, e(V_{\text{cruise}})\, d \;+\; P_{\text{hover}}\,(t_{\text{scan}} + t_{\text{winch}}) \;+\; P_{\text{compute}}\, t_{\text{scan}}, \qquad t_{\text{winch}} = \frac{2\, L_{\text{cable}}}{v_{\text{winch}}}. \tag{2}$$

Here d is the one-way distance from the launch point, doubled for the return; $e = 602.2$ J/m is the electrical energy per meter found above; $P_{\text{hover}} = 10.69$ kW; t_{scan} is the duration of the descending scan; $t_{\text{winch}} = 50$ s; and $P_{\text{compute}} = 25$ W. Hovering for the time it takes to run the cable out and back costs about as much as 0.9 km of cruise, so station-keeping dominates for a nearby site and cruise dominates near the range limit. The formula leaves out climb, acceleration, wind, diversions, and any reserve beyond the range derate.

Range decides which sites can be attempted at all. Starting from the manufacturer's published range, we scale it to the heavier analysis mass, cut it by the one-fifth battery reserve, reduce it for thinner air at 1500 m, and halve it for the return leg. The result is a one-way reach of 8.888 km. Table 11 in Section 9.1 shows the arithmetic. None of Region A's 31 blocks can be reached from the nominal regional airport, requiring 18 forward launch points to cover the region. Every reachability claim in this report depends on those forward points.

One inconsistency is recorded here and left unfixed. The range derate applies a density correction of 0.92827 for 1500 m, while the energy formula and the power curve use sea-level air. So the reachable set is pessimistic and the per-site energies are optimistic by roughly 8% for any site above sea level. Fixing this means regenerating the energy products, which we leave to future work.

7 Pre-Flight: Geospatial Site Prediction

The pre-flight stage answers one question before the aircraft leaves the ground. Using only data that exists everywhere on Earth, where in a forest could a cargo drone lower a package, and how much energy would each candidate delivery location cost to serve? The model that answers it is a segmentation network. A segmentation network labels every cell of a map, in the same way a pressure map or a noise contour map gives a value at every point.

7.1 Framework Overview

The pipeline has four stages, shown end to end in Figure 2. Harmonization stacks free satellite and terrain layers into an 18-channel raster on a 10 m grid, aligned so every layer describes the same piece of ground. Label generation, which we call the oracle, builds the training labels from airborne lidar by measuring the geometry of canopy openings directly. We had to build labels this way because no dataset of winch-delivery sites exists. Learning fits the network to those labels and calibrates its output into a

probability. Ranking pulls candidate sites out of the probability map and orders them by expected energy, using the vehicle model.

One rule organizes all of this, and we call it the input/label firewall. The model consumes only free, globally available data. High-resolution or nationally restricted data, such as airborne lidar, enters only as the source of training labels. This is what lets the model be trained where good reference data exists and used where it does not. If restricted data reached the model as an input, the model could only run where delivery is already easy. The slope gate used during label generation is the one exception, since it reads a channel that is also a model input.

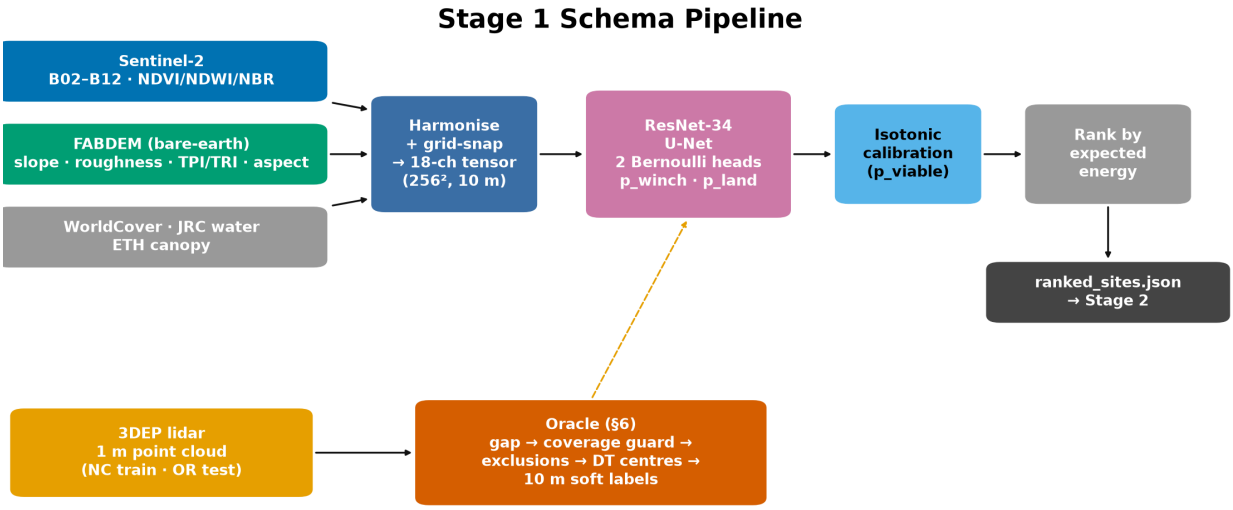

Figure 2: The pre-flight pipeline, end to end

7.2 Study Regions and Data

We needed one region to train on and one to test transfer to new terrain. Both had to have public airborne lidar of Quality Level 1 from a single snow-free campaign, complete coverage by every global input layer, and Federal ownership, meaning land held by the United States government.

Region A, used for training, validation, and in-domain testing, is the Federally owned part of the Nantahala and Pisgah National Forests in western North Carolina: 2328 km^2, with leaf-off lidar from USGS 3DEP work unit NC_Phase5_2018_A18. Restricting to Federal land removed private farmland, where the clearings are fields and yards, and would have taught the model the wrong concept. Region B is the transfer holdout, tested with no training on it at all. It is the Federally owned part of the Willamette National Forest inside the OR_McKenzieRiver_1_2021 footprint: 678.6 km^2. We picked it because its forest is the opposite of Region A's, evergreen conifer instead of deciduous hardwood, so any failure to transfer would show.

A winch center is the center point of a ground cell where a package could be lowered through the canopy on the cable. The opening has to be wide enough and the ground beneath it clear enough. Section 7.3 explains how they are measured from lidar. Raw winch centers are the ones that measurement finds, before any of the exclusions described there are applied. In Region B, 56.8% of the raw winch centers come from forest loss between 2016 and 2023, so we exclude those disturbed areas before computing any transfer result.

Table 5 lists every data product and which side of the firewall it sits on. The products are Sentinel-2 L2A [46], FABDEM [20], ESA WorldCover [47], ETH canopy height [48], USGS 3DEP lidar [49], the National Hydrography Dataset [50], OpenStreetMap [51], and USFS surface ownership [52]. FABDEM, the only free global bare-earth elevation model good enough for this work, is licensed for non-commercial use only. So the terrain data needed to serve underserved regions is free for research and not free for deployment, unless a commercial license is bought, a substitute is found, or the operation stays non-commercial.

Table 5: Data products and their role

Dataset	Product	License	Role
Sentinel-2 L2A	baselines 02.12 / 03.00	open (Copernicus)	Model input
FABDEM	v1.2, 30 m	CC BY-NC-SA 4.0	Model input
ESA WorldCover	2021 v200, 10 m	CC BY 4.0	Model input
ETH canopy height	GEDI-fused 2020 v1, 10 m	CC BY 4.0	Model input, ablation only
3DEP lidar (A)	`NC_Phase5_2018_A18`, QL1	public domain	Labels only
3DEP lidar (B)	`OR_McKenzieRiver_1_2021`, QL1	public domain	Labels only
NHD hydrography	flowline, area, waterbody	public domain	Labels only
OpenStreetMap	highway, building	ODbL	Labels only
USFS BasicOwnership	Federal surface ownership	public domain	Labels only

All layers are assembled onto one exactly specified grid from the start to avoid a misalignment between the elevation and imagery layers being mistaken for real terrain. Each 2560 m tile is a 256×256 grid at 10 m spacing. Sentinel-2 enters as a cloud-free seasonal median over leaf-on months from 2016 to 2018. Slope and the other terrain derivatives are computed at the elevation model's native 30 m spacing and only then resampled to 10 m, because resampling first stamps blocky artifacts into the derivatives that the network would learn as texture. Neighboring forest tiles look alike, so we assign whole 10 km blocks to the training, validation, and test sets with a 2.56 km buffer between them. That keeps the test blocks isolated from anything the model trained on. Region B is a separate biome about 3800 km away, as Figure 3 shows.

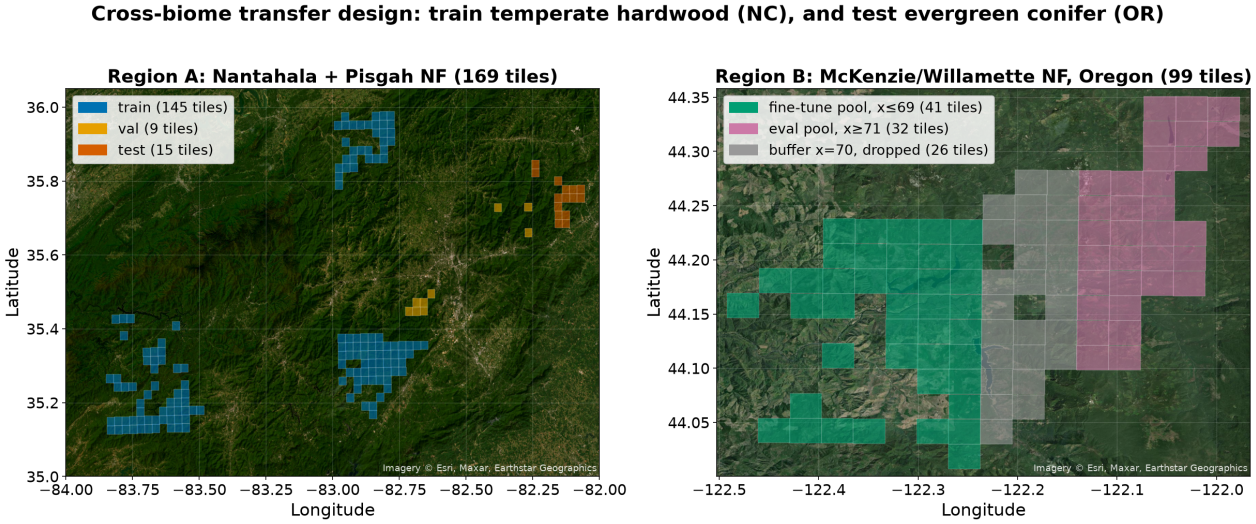

Figure 3: Spatially blocked study design

7.3 Label Generation: The Lidar Oracle

Airborne lidar is a laser scanner flown over the forest. It returns millions of 3D points wherever the laser hits a surface: a leaf, a branch, or the ground. Labels are built from those points by direct measurement and depend on terrain only: opening width, canopy height, surrounding canopy height, and slope. Cable reach, range, footprint, and payload are applied later as filters.

We call a 1 m ground cell a canopy opening if no vegetation point higher than 2.0 m above the ground lies within 0.5 m of the cell center. This is a direct query on the point cloud, with no surface built in between. A cell with no vegetation points and a cell with no lidar data at all look the same under that test, so we also check that the ground itself was densely scanned. Real clearings return 6.8 to 8.8 ground points per

m^2, while data holes return 2.4 to 4.2, and the two never overlap. That check removed 658 388 phantom openings.

The 0.5 m test radius shrinks every measured opening by that amount, so we add it back before deciding whether the 3.0 m winch footprint fits. Getting this correction wrong would demand a 7 m opening while reporting 6 m, which matters a great deal for a class of openings whose median width is 3.0 m. We checked the measurement by removing the vegetation points from disks of known size and asking the detector to find them. Openings of 8 and 10 m are measured to within 0.8 m; a 4 m opening is over-read by 1.5 to 2.0 m. So the smallest opening this method can resolve is about 6 m, and that floor bounds every claim the winch class makes. We also require the ground inside an opening to be smooth to within 0.15 m, the airframe's tolerance, which is slightly above the measurement noise of the lidar terrain model even under the densest canopy.

An opening is not automatically a place to deliver, so three exclusions follow. Water is removed using National Hydrography Dataset vectors [50] buffered by 15 m, because optical water detection misses rivers shaded by canopy; water covers 9.01% of the measured area. Slopes above 35° are removed, which drops 6.7% of candidates. The third exclusion reversed our original design. We had assumed openings in a national forest are natural. In fact 61.82% of geometrically valid winch centers lie within 50 m of an OpenStreetMap [51] road or building, while that buffer covers only 14.64% of the area. A model trained on those openings would learn to find roads, and roads are the one thing that cannot transfer to roadless terrain. Excluding them cut surviving centers from 2 462 252 to 939 993. Table 6 gives the full chain.

Table 6: Label attrition across the exclusion chain

Step	Stage	Surviving 1 m pixels	% of previous	% of raw
1	Raw canopy openings	25 493 327	—	100.00
2	Also lidar-measured	25 493 327	100.00	100.00
3	Also Federal ownership	19 730 766	77.40	77.40
4	Also away from water	16 863 461	85.47	66.15
5	Also away from roads and buildings	11 289 726	66.95	44.28
6	Also below the slope cap	10 538 640	93.35	41.34
7	Footprint fits, valid centers	939 993	8.92	3.69
7′	Same radius, roads kept	2 462 252	15.37	9.66

Step 7 is where most of the filtering happens: more than nine tenths of what survived the other exclusions is too small for a 3 m footprint. That is simply how forest is shaped. Row 7′ repeats the count with the road exclusion off, so the gap between rows 7 and 7′ is the price of refusing to learn infrastructure.

Labels are then rolled up from the 1 m grid to the 10 m model grid as a fraction: the share of each 10 m cell's measured area that is a valid winch center. Cells with no lidar coverage get no label and are left out of every metric. At 10 m spacing, a 5 m clearing cannot be resolved, so the model's output is a probability for each cell, not a landing coordinate. That is the reason the system has two stages.

We also tried the standard approach, a smoothed canopy-height surface, and rejected it on measurement: it reported zero openings at every height threshold of 3 m or more, because the smoothing fills in real clearings. We use the canopy-height surface for the cable-reach check rather than for finding gaps. The oracle is a proxy for physical viability, not ground truth, as the sites it produces have not actually been flown.

7.4 Model Architecture and Training

The model takes 18 aligned channels at 10 m: six Sentinel-2 reflectance bands; vegetation, water, and burn indices computed from them; elevation and seven terrain derivatives such as slope; and land cover class. A nineteenth channel, the ETH Zürich global canopy height model [48], is used only in the ablation experiments. One number sets the difficulty of the task: across the training tiles the vegetation index has a mean of 0.878 and a standard deviation of 0.041. That is what closed canopy looks like from orbit, and it means the model is searching for small openings in an almost constant field.

The architecture is a U-Net [53] with a ResNet-34 encoder [54], 24.48 million parameters, trained from

scratch. A U-Net compresses the image to pick up context and then expands it again, with shortcut connections that carry fine spatial detail past the compressed middle. Only the winch output is trained; the Federal-only filter left the landing class with almost no positives, by design.

The training loss combines a focal term [55] and a Tversky term [56]. In plain terms, the focal term makes the rare informative cells count for more than the sea of easy ones, and the Tversky term lets us punish misses harder than false alarms. We set it to punish a missed site 2.33 times more heavily than a false alarm, and the reason is mission cost. A missed site is never proposed and never flown to, and nothing downstream can recover it. A false alarm is a flight to a site that onboard sensing rejects, and the aircraft recovers by diverting. Recall is worth more than precision here because the in-flight stage exists to catch the false alarms.

Tiles are grouped into 10 km blocks, each assigned whole to training, validation, or test, because random splits of neighboring terrain badly underestimate error [57]. A tile is used only if at least 80% of it is Federal and it lies at least one tile width from every other split. Table 7 gives the result. Training uses the AdamW optimizer [58] at a learning rate of 3×10^{-4} for forty epochs, with no data augmentation, because flipping a tile would corrupt the direction-encoded terrain channels.

Table 7: Split composition and class balance

Split	Blocks	Tiles	Cells evaluated	Cells with a site	Share with a site
Train	20	121	7 916 642	38 347	0.004844
Validation	5	9	589 824	409	0.000693
Test	5	15	983 040	1 423	0.001448

Winch centers make up only about 0.14% of cells, so a model that predicted "no site" everywhere would be 99.86% accurate and yet useless. We therefore do not report accuracy. We report the area under the precision-recall curve, or AUPRC. Precision is the fraction of flagged cells that really are sites, recall is the fraction of real sites that were flagged, and AUPRC measures how well the model ranks true sites above false ones, which is the question a planner working down a list actually cares about. A random ranking scores an AUPRC equal to the fraction of positive cells, so we divide by that fraction and report the result as lift over chance. A lift of 38 means true sites land near the top of the list 38 times more often than by chance. Because the fraction of positive cells differs by an order of magnitude between our two regions, lift is the only number that can be compared across them, and every lift in this report is quoted with the base rate it was computed against.

A raw network output is a score, not a probability. Here the difference matters, because the ranking divides an energy by that number and the in-flight stage treats it as a belief. A score of 0.9 must mean that nine such sites in ten are deliverable. The raw model is close to calibrated on average but overconfident at the top: it emits probabilities of 1.00000 on a problem where about one cell in seven hundred is a site. We tried two standard fixes. Temperature scaling [59] rescales all outputs by one constant, and it barely moved the top (to 0.99932). Isotonic regression [60], a more flexible correction, capped the top at 0.79363 and improved the Brier score (the mean squared error of the probabilities) on both splits. We deploy isotonic regression.

We also train three copies of the model from different random starts and average their predictions, a technique called a deep ensemble [61]. Where the three copies disagree, the model is unsure, and we flag that cell. On the test split the average reaches an AUPRC of 0.08090 against a typical single copy at 0.06356 and the best single copy at 0.08687, so averaging recovers the upper end of the range but does not beat it. The disagreement flag is what the rest of the system uses; the small gain from averaging is incidental.

7.5 Site Ranking and Energy-to-Service

A probability map is not a mission plan. Candidate sites are groups of touching cells where the canopy sits well below its surroundings and the vegetation index falls below the closed-canopy value, on ground that is below 35°, not water, and Federal. Each group's position is its center and its probability is the mean calibrated model probability over the group. Two limits apply. First, the geometric rule alone decides which

groups become candidates. The learned model does not propose sites; it only scores the ones the rule has already found. Second, an opening that straddles two tiles produces two sites.

Cable reach, range, footprint, and payload are applied here, through the two winch conditions of Section 6. Sites are then ordered from lowest to highest expected total energy per successful delivery:

$$\mathbb{E}[E_{\text{total}}] = \frac{E_{\text{service}}}{\max(p_{\text{viable}},\, \epsilon_p)}, \qquad \epsilon_p = 10^{-3}. \tag{3}$$

Here E_{service} is the energy to fly to the site and service it, p_{viable} is the calibrated probability that it is deliverable, and ϵ_p is a floor that stops a near-zero probability from producing an infinite expected energy. If a site is viable with probability p, you expect to need $1/p$ attempts for one success, so expected energy is the service energy divided by the probability. This is why calibration matters, and it is the same quantity the in-flight stage optimizes one decision at a time.

Run across Region A, the pipeline produces 4404 ranked sites with calibrated probabilities from 0.0023 to 0.7936. All are reachable only from the forward launch points described with the vehicle range. One defect must be closed before operational use: the reachability flag encodes range only, and the cable-reach condition, though defined and measured, is not yet applied in the ranking code.

8 In-Flight: Deep Reinforcement Learning Agent

8.1 Framework Overview

The ranked list from the pre-flight stage is complete, but it is frozen. It was built from satellite data collected weeks or months earlier. It cannot see what the aircraft sees when it arrives, and it cannot reorder itself as the aircraft spends energy. A gap that looked open from orbit may now be closed by a fallen limb or by regrowth.

The in-flight stage handles one decision, made at one moment. The aircraft has flown to a candidate site, descended, and taken a live look at the ground. It now has to commit to a delivery here, divert to another candidate, or abort and fly home, under a hard energy budget, using the live look combined with the pre-flight probability. Each action moves the aircraft, which changes the energy cost of reaching every remaining site, and as it spends energy, it changes whether the aircraft can afford to commit to each updated energy cost. Diverting to a distant but promising site is only worth it if the battery can still cover the trip there and home. Committing to a mediocre site becomes the right call once the battery can no longer support more searching. Figure 4 illustrates the choice.

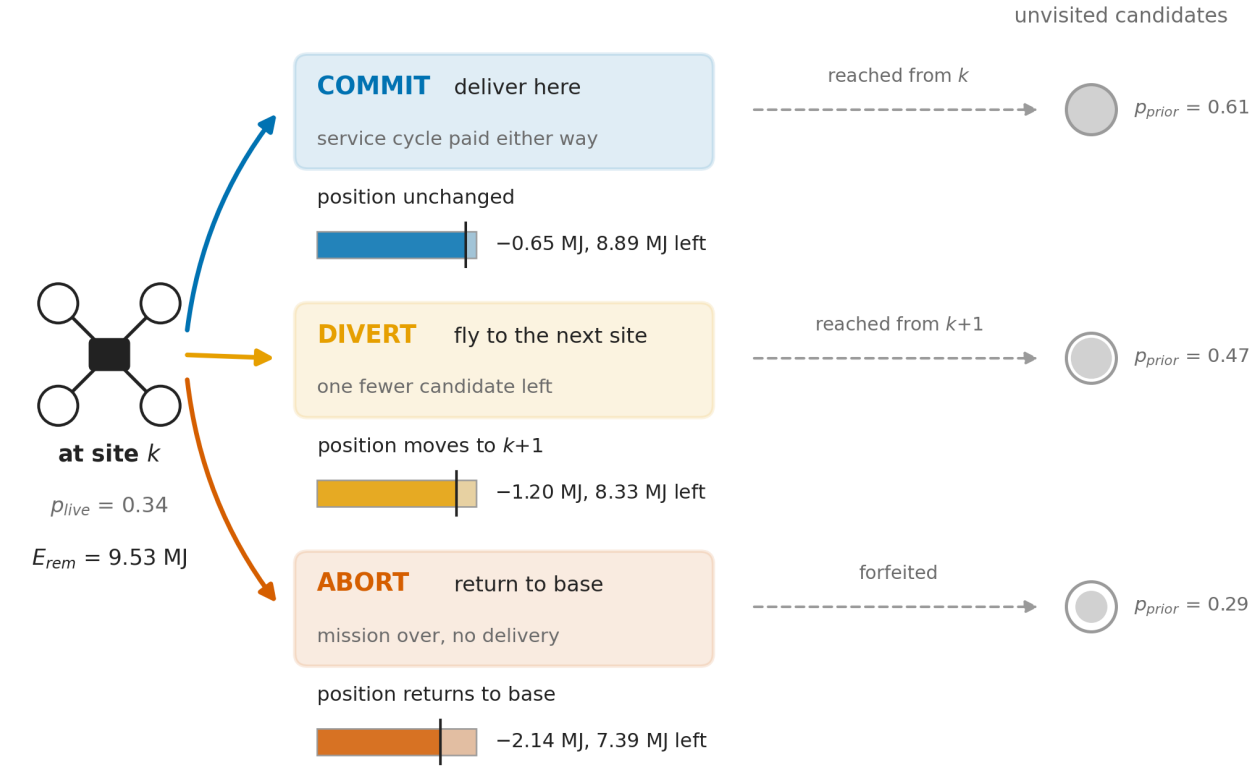

Figure 4: The decision at a candidate site

A fixed rule such as "commit if viability exceeds 0.5" cannot capture that trade, because it is blind to how much energy is left and how many candidates are still reachable. So we pose the problem as reinforcement learning, in which a policy that sees the full situation is trained over many simulated missions to minimize the same quantity the pre-flight stage ranks by: expected energy per successful delivery.

8.2 The Reinforcement Learning Framework

Reinforcement learning needs the problem stated as a state, a set of actions, a reward, and rules for how actions change the state [62]. The state is everything the aircraft knows at a decision point. The actions are commit, divert, and abort. The reward scores delivery success against energy spent and safety. What the agent learns is a value for each action in each state, meaning the total future reward it expects from taking that action and then behaving well afterwards. The best policy simply picks the highest-valued action in each state.

The true condition of a site is not part of the state. It is revealed only as a noisy observation on arrival, so the agent's estimate is a belief, and a belief may be wrong. The agent has to weigh the cost of a mistaken commit against the cost of another look. That trade is what the reward encodes.

8.2.1 Simulation Environment

A policy like this cannot be learned on a real aircraft, because learning takes on the order of 10^5 missions, including the ones that end badly. We train in a custom simulation that updates the state numerically with no graphics [63], which is what makes that many missions feasible on one processor. Each mission starts the aircraft at a base with a cluster of candidate sites in the pre-flight format. Each site carries a public probability the agent can see and a hidden condition, clear or blocked, drawn at random according to that

probability. On arrival, a sensor model produces a noisy live estimate; during training it is a statistical stand-in for the real perception module described below. Both winch and landing deliveries are represented, and the energy of every action is recomputed from the aircraft's current position at every step, because the pre-flight energies are measured from the base and are only correct for the first leg out.

8.2.2 What the Agent Sees and What It Can Do

The state is a fixed list of 36 numbers. Nine describe the aircraft and the site it is over, and 27 describe the next nine unvisited candidates in rank order, three values each. Table 8 gives the composition. Unused slots are filled with zeros.

Table 8: Composition of the state vector

Values	Quantity	Role
3	Position east, north, and up, as an offset from base	Sets the marginal energy to every remaining candidate
1	Remaining energy as a fraction of the usable budget	Whether a given action is affordable at all
1	Fraction of the energy budget already spent	Pacing: how decisively to act late in a mission
1	Index of the current site in the ranked list	Position within the list
2	Live viability estimate and its confidence	The look the aircraft has just taken
1	Fused viability estimate	Live estimate combined with the pre-flight prior
27	For each of the nine next candidates: pre-flight probability, extra energy to fly there, delivery mode	What each remaining option is worth and what it costs from here

Remaining energy and the fraction of the budget already spent are both included on purpose: the first says whether an action is affordable, the second says how far through the mission the aircraft is. The usable budget is the installed battery minus the one-fifth reserve, the same reserve the range derate uses. Nine candidates is the ninetieth percentile of how many sites an aircraft can still reach within its budget. The hidden condition of each site is never in the state, because an agent that could see it could not exist in deployment.

Commit attempts a delivery here, abort returns to base, and divert flies to the next candidate. The divert target is fixed by rule as the nearest reachable unvisited site by energy, so the agent chooses whether to divert, not where. Letting it choose the target would grow the action space from three to more than a dozen and slow learning, and flying to the cheapest reachable site is not a decision that benefits from learning. An action mask blocks illegal moves before the agent picks: commit is blocked at a site infeasible for its delivery mode and at a flagged site until a live look has been taken, divert is blocked when the trip there and home would breach the reserve, and abort is always allowed.

8.2.3 Reward Function

The reward states the objective, and it has to say four things at once: a completed delivery is the goal, energy spent is a cost, finishing below the safety reserve is worse than failing, and running the battery dry is unacceptable at any price.

$$\begin{aligned} r_t = R_{\text{success}}\,\mathbb{1}[\text{delivered}] \;-\; w_{\text{energy}}\,\Delta E_t \;-\; C_{\text{abort}}\,\mathbb{1}[\text{abort}] \\ -\; C_{\text{reserve}}\,\mathbb{1}[\text{reserve breach}] \;-\; C_{\text{crash}}\,\mathbb{1}[\text{exhaustion}], \end{aligned} \tag{4}$$

$$w_{\text{energy}} = R_{\text{success}} \,/\, E_{\text{usable}}.$$

Each indicator fires only when the named event happens on that step, ΔE_t is the energy spent by the action, and E_{usable} is the usable battery budget. Table 9 gives the constants.

Table 9: Reward constants and their basis

Term	Value	Basis
Reward for a delivery	1	Sets the scale; every other term is referred to it
Energy penalty weight	Delivery reward divided by usable energy	Derived, so spending a full battery costs the same as one delivery earns
Penalty for battery exhaustion	Far above the delivery reward	No sequence of deliveries an episode can reach may pay for exhaustion
Penalty for breaching the reserve	Above the delivery reward	A single delivery may not pay for eating into the safety margin
Penalty for aborting	Far below the delivery reward	Small; discourages abandoning a recoverable mission
Penalty for a failed commit	None	Already charged through the energy term

A delivery is worth one unit, and the energy weight follows from that, so spending the entire usable battery costs exactly one unit. That keeps the constant valid if battery capacity changes. There is no separate penalty for a failed commit, because a failed commit has already paid for its energy, and charging it twice would make the policy too timid. This reward defines a break-even viability, the probability at which committing and diverting are worth the same. That threshold depends on how much an attempt costs: where a failed attempt is expensive the threshold is high, and where it is cheap the threshold falls and trial-and-error becomes efficient.

8.2.4 Learning Algorithm

The action values are learned by a neural network that maps the 36-number state to three values. It is trained toward a self-consistent target: the value of an action should equal the reward received plus the value of the state that follows [64]. The standard version of this method uses one network both to pick the best next action and to score it, which biases it toward whichever action happens to be overestimated by noise [65]. Here that bias is dangerous, because an overestimated commit or divert is an optimistic belief that an action is affordable when it may strand the aircraft. So we use the double deep Q-network, in which one network picks and a separate, slowly updated copy scores. Masked actions are excluded from the target, so the agent never learns from an illegal move. Exploration takes a random action with a probability that decays over training. The crash penalty produces occasional very large errors, so we train with the Huber loss [66], which caps how much any one transition can move the network.

8.2.5 Onboard Perception and Sensor Fusion

The decision agent uses a live viability estimate but does not produce one. That comes from a separate perception module, trained offline and frozen at decision time, which maps the onboard sensor data at a site to a probability and a confidence. It sees neither the energy budget nor the candidate list.

Three sensors are fused: the visible branch of the added FLIR camera, its thermal branch, and the airframe's factory lidar depth sensor. They fail in different ways, which is the point. A gap that looks ambiguous in shadow to the visible camera is plain to the depth sensor. Each sensor is processed by its own convolutional branch, and the three summaries are combined only just before the output, as Figure 5 shows. Combining them this late means a degraded sensor does not take down the others, and each branch can be tested alone.

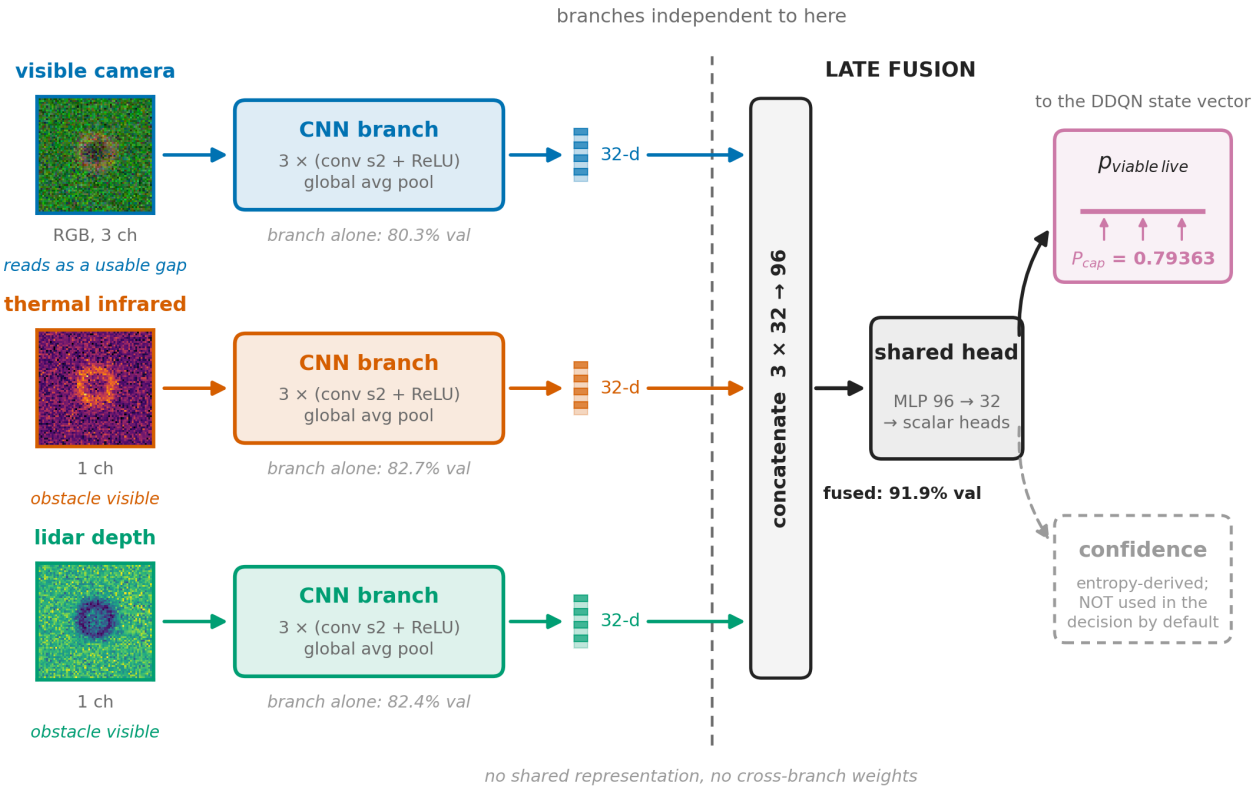

Figure 5: Perception module

The module's probability is capped at 0.79363, the same ceiling the pre-flight calibration produced, and for the same reason: a raw network is overconfident at the top, and that is where commitment decisions are made. That two different models on two different data sets landed on the same ceiling suggests a real limit on how certain this kind of prediction can be. Using one cap in both stages keeps a probability meaning the same thing on both sides of the interface. The confidence output correlates only weakly with correctness, so it is kept out of the decision by default.

9 Results

Two conventions apply throughout. Every lift ratio is quoted with the base rate it was computed against, because base rates differ by nearly nineteen times between the two regions. And every ± spread is a standard deviation across training runs started from different random seeds. It measures how much the result varies by chance, not how well the model generalizes to new terrain.

9.1 Vehicle and Energy Results

The cable-length requirement is the result we trust most, because it compares the canopy of a real national forest against a published cable length and no model enters it. Measured against the tallest tree around each site, which is the tree the aircraft has to clear, the installed 30 m cable reaches 42.69% of arbitrary forest and 52.15% of deliverable sites. A 40 m cable would reach 91.52% and 90.64%. Table 10 sweeps the length, and the curve bends sharply between those two values: each extra 5 m buys a lot below

40 m and very little above it. The aircraft as sold serves about half the sites this work finds. The length needed to serve nine in ten is a requirement this analysis produces.

Table 10: Fraction served against installed cable length

Cable length (m)	20	25	30	35	40	45	50
Arbitrary forest reached (%)	9.53	21.11	42.69	68.59	91.52	98.27	99.82
Deliverable sites reached (%)	—	37.42	52.15	69.79	90.64	—	—

The two rows are measured over different populations, 14 302 627 forest cells for the first and 652 proposed sites for the second, so compare them down a column, not across. Read that way, sites are only about a fifth easier to serve than arbitrary forest. The advantage is small because sites are already selected for sitting in openings: the tallest surrounding tree has a median of 23.33 m against an area-wide canopy mean of 31.4 m, and its ninetieth percentile of 34.75 m plus 5 m of clearance is 39.75 m, which is where the 40 m figure comes from. Although we found that 40 m would expand the set of feasible sites, because the vehicle energy model assumed factory installed 30 m winch length, all remaining analyses still assume 30 m. Extending the research to the 40 m winch case is a consideration for future work.

No part of Region A can be served from a single origin. At the derated one-way range of 8.888 km, the nearest block is 10.280 km from the nominal regional airport and the median block is 54.364 km. Covering the 31 usable blocks needs 18 forward launch points; that set achieves coverage but was not proven minimal. Table 11 gives the derate chain. Only the fourth row involves physics: air at 1500 m is thinner, so the rotors move more of it for the same lift. Usable energy after the reserve is 12.28 MJ of an installed 15.3504 MJ.

Table 11: Range derate chain

Step	Adjustment	Range (km)
Published range at reference condition	—	12.000
Scaled to the 94.6 kg analysis mass	weight to the power −1.5	23.936
One-fifth energy reserve held back	× 0.80	19.149
Thinner air at 1500 m	× 0.928	17.775
Halved for the return leg	× 0.50	8.888

One caveat applies to every absolute energy figure. The pre-flight and in-flight codebases disagree on electrical hover power by roughly 45%, and neither contains a forward-flight power curve; the pre-flight site energies use an undocumented ratio of 0.7 times hover power at the published route speed. Absolute energies should therefore be read as order-of-magnitude quantities. The relative comparisons below each change one thing at a time inside one codebase, and are unaffected.

9.2 Pre-Flight Model Performance

The case for a learned model starts with a negative result. The global canopy height product the model sees at inference is nearly blind to the openings that define the winch class. Over lidar-confirmed openings it reports a median canopy height of 25.0 m, against 31.0 m over closed canopy, and it shows at least 10 m of canopy over 97.6% of confirmed open cells. A simple threshold on that product finds 41 winchable cells across the entire study area. That number is the reason to use learning at all. Ranking those geometric candidates by expected energy instead of at random cuts mean service energy from 6.1787 MJ to 1.8457 MJ, a factor of 3.35, though this is a reordering of a fixed list, not a flown mission.

The learned model reaches a test AUPRC of 0.05568 ± 0.01225 over ten seeds, a lift of 38.5 ± 8.5 at a base rate of 0.001448. The seed spread is 22% of the mean, so no single run means anything on its own. We record a retraction on that basis: an earlier headline of 0.07013 was a real run but an unrepresentative one, sitting well above a mean whose runs span 0.032 to 0.071.

Table 12: Channel ablation ladder

Configuration	Channels	Seeds	Test AUPRC	Lift over chance
RGB imagery only	3	10	0.04146 ± 0.01502	28.6
Plus band indices	6	5	0.04627 ± 0.01637	32.0
Plus slope	9	5	0.04859 ± 0.01285	33.6
Plus canopy height	10	5	0.04996 ± 0.00885	34.5
Full stack less canopy height	18	5	0.05152 ± 0.01298	35.6
Full stack less slope	18	5	0.05251 ± 0.00819	36.3
Full stack	19	10	0.05568 ± 0.01225	38.5

Table 12 adds inputs one group at a time, from the plain red, green, and blue bands up to the full stack, all at a base rate of 0.001448. The mean rises at every step, but every step is inside one standard deviation, so only the endpoints can be compared and even that is marginal: the full stack gains 34.3% over imagery alone with seven of ten seeds positive and a 95% interval that includes zero by 0.0006. More seeds would not fix this, because all ten are scored on the same five test blocks, whose gap density varies tenfold. The fix is more held-out blocks, which we did not have.

The canopy height channel contributes nothing measurable, and two independent tests agree. Adding it to the nine-channel stack gains 0.0014 ($p = 0.7761$); removing it from the full stack costs 0.0062, well inside one standard deviation. This is the gap-blindness result again: a channel that cannot see openings cannot help find them.

Calibration behaves as intended. Isotonic regression wins the Brier score on both splits and is the only method that improves test calibration error over the raw model. The mean predicted probability on true sites is 0.0593, so the model is strongly under-confident on the positive class, which is what the recall-first loss was meant to do.

Deferral is the result the architecture was built around. A cell on which the three model copies disagree is flagged as uncertain, and that flag can be spent two ways: keep the site on the list and check it from the air, or drop it before launch. Table 13 compares both against deferring nothing, with the most uncertain 5% of cells flagged.

Table 13: Effect of deferring the most uncertain cells

Policy	Expected energy (MJ)	Median energy (MJ)	Mission success
No deferral	77.593	10.061	0.60
Defer to onboard verification	50.138	10.214	0.60
Defer by discarding	8.559	—	0.067

Checking uncertain sites from the air cuts expected energy per successful delivery by 35.4% at no cost in recall. Dropping them cuts energy much further but collapses mission success by an order of magnitude, because the uncertain cells are where the real sites are: the flagged 5% has a base rate of 0.00981 against 0.001448 for the whole population, 6.8 times higher. Uncertainty must mean verify, not discard. The gain is all in the tail. The mean falls 35% while the median barely changes, so deferral removes catastrophic missions and leaves the typical one alone. These are offline expectations over the ranked list, not flown missions.

Two limits bound the winch class itself. Only 13 of 4194 measured openings, 0.31%, are 6 m or wider, against a most common width of 3.0 m. And the class is very sensitive to footprint radius: the count of valid centers falls from 1302 to 113 as the radius moves from 2.5 m to 4.0 m, so half a meter of assumed clearance costs about two thirds of the class.

9.3 Cross-Biome Transfer and Label Efficiency

Region B has a base rate of 0.028 against Region A's 0.001448, roughly nineteen times higher. The difference is physical: Pacific Northwest conifer forest is far more open than Appalachian hardwood, so more

of it really is winchable. Lift is therefore the transfer metric.

Table 14: Configuration ranking reverses between regions

Configuration	Channels	Region A lift	Region B lift
RGB imagery only	3	28.6	2.85 ± 1.13
Plus band indices	6	32.0	2.92 ± 1.16
Plus slope	9	33.6	3.02 ± 1.02
Plus canopy height	10	34.5	2.80 ± 1.11
Full stack less canopy height	18	35.6	2.42 ± 0.75
Full stack less slope	18	36.3	2.23 ± 0.74
Full stack	19	38.5	2.01 ± 0.67

Table 14 puts each input configuration's in-domain lift next to its lift in Region B with no training there. The ranking reverses. In-domain lift rises down the rows while out-of-domain lift falls, so the configuration that is best in Region A is worst in Region B, and the modest nine-channel stack transfers best. Every gap between rows is inside one standard deviation; the strength of the result is that all seven rows reverse together, which is far harder to produce by chance than any single comparison. The mechanism is visible in the inputs: between regions the satellite reflectance channels shift by 1.1 to 3.9 Region A standard deviations and canopy height by 2.7, while the terrain channels move by at most 0.62. The more a configuration leans on the shifted channels, the further out of its training range it is pushed.

Splitting Region B into natural gaps and disturbance openings such as burn scars sharpens the picture, because Region A's labels are natural gaps by construction and about a quarter of Region B is recent forest loss [67]. On natural gaps every configuration reaches roughly 1.7 to 2.0 times its lift on disturbance openings, where the full stack barely beats chance at 1.18. The reversal appears in both strata.

Table 15: Label efficiency on the natural stratum

Fine-tuning tiles	Plus slope (9 ch)	Full less canopy (18 ch)	Full stack (19 ch)
0, no fine-tuning	3.02 ± 1.02	2.42 ± 0.75	2.01 ± 0.67
3	4.09 ± 1.09	3.76 ± 0.69	4.28 ± 0.94
6	4.58 ± 0.97	4.31 ± 0.49	5.10 ± 1.08
12	4.97 ± 0.67	5.47 ± 1.13	5.59 ± 0.83
24	5.23 ± 0.57	5.61 ± 0.77	5.59 ± 0.72
41, all available	5.21 ± 0.53	5.92 ± 0.57	5.80 ± 0.53

Table 15 asks the practical question: how many labeled tiles from the new region does it take to recover? Three tiles recover most of the loss. For the full stack, lift moves from 2.01 with no local training to 4.28 with three tiles, 60% of the way to the ceiling, and 82% of the way by six. Saturation arrives near twelve: from twelve to forty-one tiles, a 3.4-fold increase in labels, the full stack moves only from 5.59 to 5.80. Figure 6 shows the arc. The full stack is best in domain, worst with no local training, and best again after three local tiles. So the transfer penalty is a zero-shot-only problem that a handful of local labels dissolves, and entering a new region costs about a dozen labeled tiles, not a national lidar campaign. Only three of the seven configurations were run for this experiment.

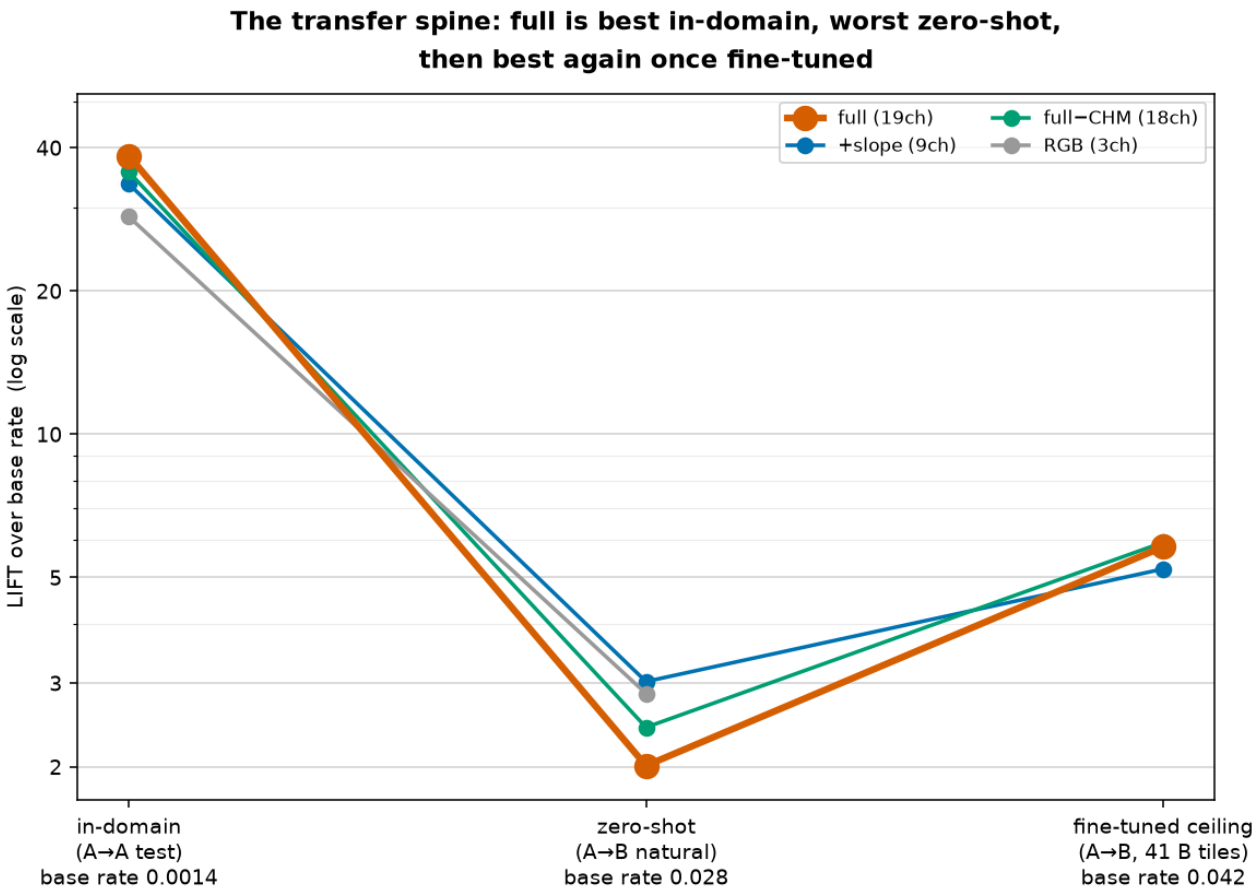

Figure 6: The transfer arc across three regimes

9.4 In-Flight Policy and Mission Performance

The agent was trained over 40 000 simulated missions, with reward plateauing near mission 20 000 and the best checkpoint taken at 15 000, on a single seed. Every figure below is evaluated on a fixed set of 1000 held-out maps, with intervals computed paired per map over the 894 maps where a delivery was possible.

The learned policy commits on viability and aborts on energy. Across increasing bins of live viability, the commit rate rises from 6.1% to 58.1% while the divert rate falls from 89.4% to 41.2%. Abort fires at 21.9% in the lowest budget bin and at 0.8% or below in every higher bin. Commitment plotted against remaining energy is roughly flat in Figure 7, so the policy does not gamble more on doubtful sites as the battery runs down.

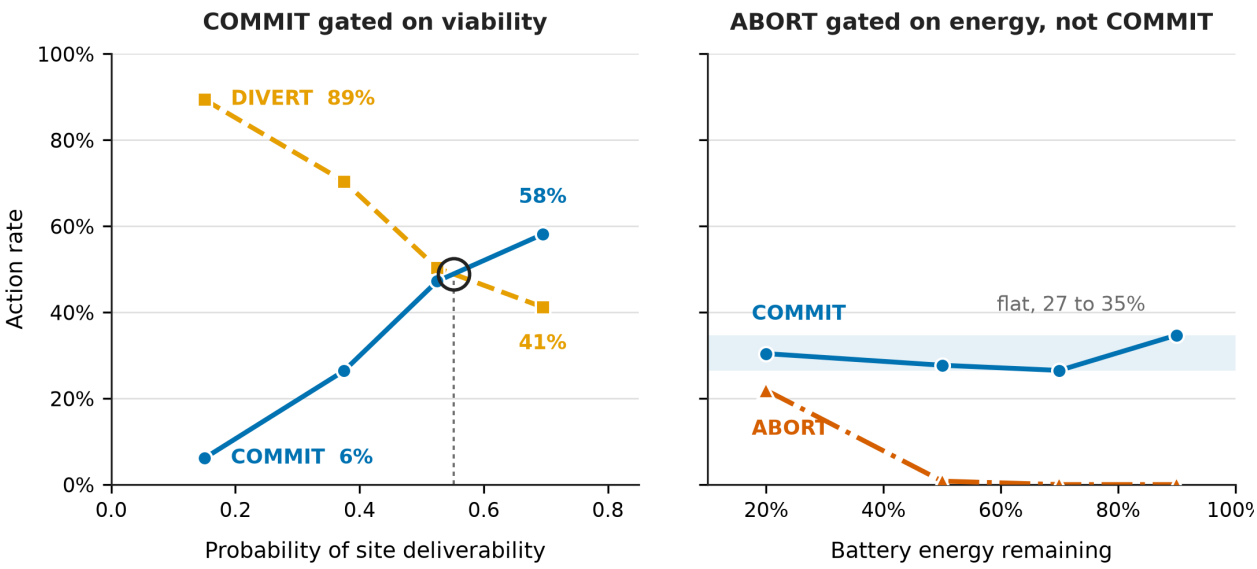

Figure 7: Action rates of the learned policy

Safety needs a correction to the claim made with the action mask. True battery exhaustion is 0.0% on the evaluation set, but the reserve is breached in 1.1% of missions. The divert guard prevents breach only for a policy that only diverts, and commit carries no energy guard, so crash-free operation is a measured outcome, not a guarantee. Closing that gap is a design change.

Table 16: Single modality against late fusion

Modality	Overall accuracy (%)	Blocked sites rejected (%)	Clear sites recognized (%)
Visible	80.3	86.8	73.1
Thermal	82.7	88.0	77.0
Lidar depth	82.4	90.6	73.4
Late fusion	91.9	89.9	94.0

Table 16 compares each sensor alone against the fused module on a balanced validation set of 800 samples. Fusion gains 9.2 points of accuracy over the best single sensor and raises clear recall from 73.4% to 94.0% while blocked rejection moves only -0.7 points. It recovers the good sites the single sensors were throwing away without accepting more bad ones. At the deployment base rate, over 2500 held-out predictions with 30.9% of sites clear, the module rejects 89.8% of truly blocked sites (95% interval $[88.2, 91.1]$) and recognizes 93.5% of clear ones. That rejection rate is now measured; the pre-flight false-positive rate is not, so the recall-first argument is closed on one side only.

Table 17: Learned policy against controls

Comparison	Full agent	Control	Paired difference, 95% interval
Delivery rate against a control that cannot see the pre-flight prior	96.1%	93.7%	+2.3 points, $[+0.9, +3.8]$
Delivery rate against a control that diverts by rank order rather than by energy	96.1%	91.8%	+4.3 points, $[+2.3, +6.2]$

Table 17 compares the learned policy against two controls on the same 1000 maps. Seeing the pre-flight probability is worth 2.3 points of delivery rate. Choosing the divert target by energy instead of by rank order is worth nearly twice as much. The prior-blind control still diverts by nearest energy, so it is a strong control. In simulation, inspecting before committing saves 5.8% of mean energy per delivery and avoids 930 false winch cycles, against the 35.4% predicted offline; the offline figure assumed a perfect verifier, so the gap is the price of an imperfect look.

The corrected vehicle physics changed what the agent is for. As service energy fell from 2.25 MJ to 0.645 MJ, the break-even viability fell from about 0.17 to 0.061, below the typical blocked-site probability of 0.128. Attempts became cheap enough that stopping early no longer saves much energy. The reason to inspect before committing has therefore moved from saving energy to protecting the reserve and the delivery rate.

One disclosure applies to everything above. The two stages have been checked against each other's data formats, probability ceiling, coordinate frame, and energy model, but they have never been run as one system. The agent was trained and evaluated on synthetic ten-site clusters shaped to the interface, not on the 4404-site regional product, with probabilities not matched to that product and synthetic perception inputs in place of real imagery. The interface is validated. The coupling is not. The layer that reduces a region to one mission's candidates is the remaining integration step.

10 Conclusions

This work set out to choose delivery sites in terrain that has never been surveyed, and it split the task in two. A pre-flight stage computes a probability map over a whole region from satellite and elevation data that is free and available everywhere, trained against airborne lidar that is neither free to use commercially nor available outside a few wealthy countries. An in-flight stage carries that map aloft, verifies candidates during descent, and diverts when one fails, where a conventional system would abort. Both stages minimize expected energy per successful delivery.

The result we trust most is a vehicle requirement. Measured against the canopy of an actual national forest, the installed 30 m cable reaches roughly half of deliverable sites, where 40 m would reach nine in ten. No model, training seed, or threshold enters that number. The second result is the shape of transfer between forest types: the configuration that is best at home is the worst of seven abroad, three local tiles recover 60% of the loss and twelve saturate it, and the cause is visible in the inputs, where spectral channels shift far outside their home range while terrain channels barely move. Entering a new biome costs about a dozen labeled tiles, not a national survey. The third result is that uncertainty must mean verify, not discard, because the sites the model is unsure about are disproportionately the real ones.

The clearest operational finding is negative. No block of the study region is reachable from the regional airport, and covering it needs eighteen forward launch points, so an aircraft of this class cannot serve a region of this size from one depot. Several other negative results carry weight: the global canopy product is nearly blind to the openings that matter, so a simple threshold finds 41 sites in an entire national forest; adding a canopy channel buys nothing measurable; one of three pre-registered transfer predictions failed in the opposite direction; and cheap service energy moved the case for in-flight inspection from energy to safety.

Four conditions bound all of this. Every result assumes the installed 30 m cable; 40 m is a requirement we derive, not a configuration we ran. The two codebases disagree on hover power by about 45%, so absolute energies are order-of-magnitude while relative comparisons hold. Every reported interval measures seed-to-seed noise, not spatial generalization. And the two stages are assumed to compose, because they have never been run together, the pre-flight stage has never measured its false-positive rate, and transfer was measured in one direction between one pair of biomes.

11 Future Work

What this system needs next is measurement more than new capability. The highest-value item is short: no operating threshold is selected anywhere in the pre-flight pipeline, so the false-positive rate that the

recall-first design depends on has never been measured, and evaluating one on the held-out tiles is a quick run. The energy model needs a forward-flight power curve inside both codebases, which would replace the undocumented cruise ratio, reconcile the two hover powers, and make headwind answerable. The products also need regenerating for the 40 m cable, the length this analysis identifies as the one that would reach nine deliverable sites in ten. The layer that reduces a region-wide site list to one mission's candidates does not yet exist, and building it is the step that makes the two stages one system.

Several results need stronger evidence: a block-level bootstrap over the test blocks would give the first interval that speaks to spatial generalization, the in-flight policy was trained on one seed, and the opening-width census needs re-running with the resolution correction applied. The transfer study should extend beyond one pair of biomes, and if spaceborne lidar can replace airborne lidar in the label generator, the twelve-tile budget could be met with satellite passes instead of survey flights. The winch pendulum derivation was never written, so the footprint radius that defines the whole label set rests on an assumed swing allowance. The perception module has been evaluated only on synthetic imagery. Finally, this work asks only where a package can be placed, not where one is needed; recipient and demand modeling would close the other half of the problem.

Appendix

Reference Vehicle Parameters

Table 18 collects every parameter the power and mission-energy models use, with the origin of each. Manufacturer entries come from the published specification sheet. Computed entries follow arithmetically from entries above them. Derived entries are worked out from independent data. Assumed entries are engineering allowances we set and did not verify by measurement, and they are the entries a follow-on study should test first.

Table 18: Reference vehicle parameters and their sources

Parameter	Symbol	Value	Source
Rotor radius	R	0.7874 m	Manufacturer [18]
Area of one rotor disc	A_{rotor}	1.948 m^2	Computed, πR^2
Blade-swept area, profile power	A_{pro}	15.584 m^2	Computed, eight rotor discs
Effective disc area, induced power	A_{ind}	7.792 m^2	Computed, four shared discs
Rotor angular speed	Ω	90 rad/s	Derived from propulsion bench data [43]
Blade tip speed	V_{tip}	70.87 m/s	Computed, ΩR
Airframe and winch mass	—	60.2 kg	Manufacturer [18]
Battery mass, two packs	—	29.4 kg	Manufacturer [18]
Payload mass	—	5.0 kg	Configuration parameter
Analysis mass	m	94.6 kg	Computed, sum of the three above
Weight at analysis mass	mg	928.0 N	Computed
Installed battery energy	E_{batt}	15.3504 MJ	Manufacturer, 4264 Wh [18]
Cable length	L_{cable}	30 m	Manufacturer [18]
Winch line speed	v_{winch}	1.2 m/s	Manufacturer [18]
Clear radius for a winch delivery	r_{winch}	3.0 m	Assumed
Clear radius for a landing	r_{land}	4.0 m	Assumed
Equivalent flat-plate area	f_{par}	1.44 m^2	Derived by bluff-body analogy [18, 35, 36, 37, 38, 39]
Induced-power correction factor	κ	1.28	Coaxial interference model [45]
Solidity–drag product	σC_{d0}	8.7×10^{-4}	Derived from bench data [43, 34, 44]
Profile-power advance-ratio factor	k_μ	4.6	Standard rotorcraft value [34]
Propulsion chain efficiency	η_{prop}	0.85	Assumed; motor, controller, and propeller losses
Electrical hover power	P_{hover}	10.69 kW	Derived: published hover endurance [18] scaled to the analysis mass
Avionics bus load	P_{avionics}	78.5 W	Assumed allowance
Air density	ρ	1.225 kg/m^3	Standard sea level
Airspeed envelope for the power curve	$V_{\max}$	20 m/s	Sweep bound [18]

References

[1] Kruk, M. E.; Gage, A. D.; Joseph, N. T.; Danaei, G.; García-Saisó, S.; and Salomon, J. A.: Mortality due to low-quality health systems in the universal health coverage era: a systematic analysis of amenable deaths in 137 countries. The Lancet, vol. 392, no. 10160, 2018, pp. 2203–2212. doi:10.1016/S0140-6736(18)31668-4.

[2] Scheil-Adlung, X., ed.: Global Evidence on Inequities in Rural Health Protection: New Data on Rural Deficits in Health Coverage for 174 Countries. ILO Extension of Social Security (ESS) Paper Series No. 47, International Labour Office, Geneva, 27 April 2015, 83 pp. https://www.social-protection.org/gimi/gess/Media.action?id=14833

[3] Knight, H. E.; Self, A.; and Kennedy, S. H.: Why are women dying when they reach hospital on time? A systematic review of the 'third delay'. PLOS ONE, vol. 8, no. 5, 21 May 2013, e63846. doi:10.1371/journal.pone.0063846.

[4] Vledder, M.; Friedman, J.; Sjöblom, M.; Brown, T.; and Yadav, P.: Improving supply chain for essential drugs in low-income countries: results from a large scale randomized experiment in Zambia. Health Systems and Reform, vol. 5, no. 2, 2019, pp. 158–177. doi:10.1080/23288604.2019.1596050.

[5] Romeo, R.; Vita, A.; Testolin, R.; and Hofer, T.: Mapping the Vulnerability of Mountain Peoples to Food Insecurity. Food and Agriculture Organization of the United Nations, Rome, 2015. https://openknowledge.fao.org/server/api/core/bitstreams/fc51a31f-4d11-45da-a9f3-5d44277ab231/content

[6] Human populations in the world's mountains: spatio-temporal patterns and potential controls. PLOS ONE, 2022. doi:10.1371/journal.pone.0271466.

[7] Newton, P.; Kinzer, A. T.; Miller, D. C.; Oldekop, J. A.; and Agrawal, A.: The number and spatial distribution of forest-proximate people globally. One Earth, vol. 3, no. 3, 18 September 2020, pp. 363–370. doi:10.1016/j.oneear.2020.08.016.

[8] Zipline surpasses 2 million deliveries with expansion to Houston and Phoenix. DroneLife, 21 January 2026. https://dronelife.com/2026/01/21/zipline-surpasses-2-million-deliveries-with-expansion-to-houston-and-phoenix/

[9] Cleveland Clinic Launches Drone Delivery Program for Prescription Medications. Cleveland Clinic Newsroom, 3 August 2026. https://newsroom.clevelandclinic.org/2026/08/03/cleveland-clinic-launches-drone-delivery-program-for-prescription-medications

[10] Guarino, B.: How Zipline designed its droid delivery system. IEEE Spectrum, 2024. https://spectrum.ieee.org/delivery-drone-zipline-design

[11] Zipline adds rappelling droid to delivery drones. IEEE Spectrum, 2024. https://spectrum.ieee.org/zipline-drone-delivery-2659602543

[12] Zipline is expanding its autonomous drone delivery services to the Phoenix area. Axios Phoenix, 9 April 2026. https://www.axios.com/local/phoenix/2026/04/09/zipline-autonomous-drone-deliveries-valley-expansion

[13] Aggarwal, S.; Gupta, P.; Mahajan, N.; Balaji, S.; Singh, K. J.; Bhargava, B.; and Panda, S.: Implementation of drone based delivery of medical supplies in North-East India: experiences, challenges and adopted strategies. Frontiers in Public Health, vol. 11, 2 June 2023, art. 1128886. doi:10.3389/fpubh.2023.1128886.

[14] Drone network reduces medical delivery time in Hainan's mountainous regions. People's Daily Online, 14 May 2026. https://en.people.cn/n3/2026/0514/c98649-20456126.html

[15] First-ever cargo drone deliveries in the Amazon rainforest. WeRobotics, 2017. https://werobotics.org/blog/first-ever-cargo-drone-deliveries-in-amazon-rainforest-2

[16] Mohammed, M. M.; El Gelany, S.; Eladwy, A. R.; et al.: A ten year analysis of maternal deaths in a tertiary hospital using the three delays model. BMC Pregnancy and Childbirth, vol. 20, no. 1, 6 October 2020, art. 585. doi:10.1186/s12884-020-03262-7.

[17] U.S. Government Accountability Office: Air Ambulance — Available Data Show Privately-Insured Patients Are at Financial Risk. GAO-19-292, 20 March 2019. https://www.gao.gov/assets/gao-19-292.pdf

[18] DJI: FlyCart 100 Specifications. https://www.dji.com/flycart-100/specs

[19] Stierlin, N.; Risch, M.; and Risch, L.: Current advancements in drone technology for medical sample transportation. Logistics, vol. 8, no. 4, 12 October 2024, art. 104. doi:10.3390/logistics8040104.

[20] Hawker, L.; Uhe, P.; Paulo, L.; Sosa, J.; Savage, J.; Sampson, C.; and Neal, J.: A 30 m global map of elevation with forests and buildings removed. Environmental Research Letters, vol. 17, no. 2, 3 February 2022, art. 024016. doi:10.1088/1748-9326/ac4d4f.

[21] Leitold, V.; Keller, M.; Morton, D. C.; Cook, B. D.; and Shimabukuro, Y. E.: Airborne lidar-based estimates of tropical forest structure in complex terrain: opportunities and trade-offs for REDD+. Carbon Balance and Management, vol. 10, no. 1, 3 February 2015, art. 3. doi:10.1186/s13021-015-0013-x.

[22] Federal Aviation Administration: Small Unmanned Aircraft Systems. 14 CFR Part 107; Docket FAA-2015-0150, 81 FR 42209, 28 June 2016, as amended. https://www.ecfr.gov/current/title-14/chapter-I/subchapter-F/part-107

[23] Federal Aviation Administration and Transportation Security Administration: Normalizing Unmanned Aircraft Systems Operations Beyond Visual Line of Sight. Notice of Proposed Rulemaking, 90 FR 38212, Docket No. FAA-2025-1908, RIN 2120-AL82, 7 August 2025. https://www.federalregister.gov/documents/2025/08/07/2025-14992/

[24] Executive Order 14307, Unleashing American Drone Dominance, 6 June 2025; 90 FR 24727–24731, 11 June 2025. https://www.federalregister.gov/documents/2025/06/11/2025-10814/

[25] FAA Reauthorization Act of 2024, Section 930. Public Law 118-63, adding 49 U.S.C. 44811.

[26] Federal Aviation Administration: Written Re-Evaluation and Record of Decision, Amazon Prime Air Amendment to Operations Specifications, Pendleton, Oregon. 27 August 2025. https://www.faa.gov/uas/advanced_operations/nepa_and_drones/20250827_Amazon_Pendleton_OR_Written_ReEvaluation.pdf

[27] Manufacturer specification sheets for the candidate aircraft of Table 3. DJI, Wingcopter, Draganfly, Harris Aerial, Vision Aerial (https://visionaerial.com/vector/), Quantum Systems, Skydio, ModalAI, Freefly, and A2Z Drone Delivery. [all accessed 7 August 2026]

[28] American Security Drone Act of 2023. Title XVIII Subtitle B, Sections 1821–1833 of the National Defense Authorization Act for Fiscal Year 2024, Public Law 118-31, 22 December 2023; 41 U.S.C. 3901 note prec.

[29] Federal Communications Commission: Covered List — addition of foreign-produced unmanned aircraft systems and UAS critical components. Effective 23 December 2025. https://www.fcc.gov/document/fcc-updates-covered-list-add-certain-uas-and-uas-components-0

[30] National Defense Authorization Act for Fiscal Year 2025, Section 1709 (Countering CCP Drones Act).

[31] NVIDIA: Jetson Orin Nano Super Developer Kit specifications. https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/nano-super-developer-kit/

[32] Teledyne FLIR: Hadron 640 Series — High Performance, Dual Thermal and Visible OEM Camera Module. https://oem.flir.com/products/hadron-640

[33] DJI: Payload SDK developer documentation. https://developer.dji.com/doc/payload-sdk-api-reference/v2/en/

[34] Leishman, J. G.: Principles of Helicopter Aerodynamics. 2nd ed., Cambridge University Press, Cambridge, 2006.

[35] Hoerner, S. F.: Fluid-Dynamic Drag. Published by the author, Bakersfield, CA, 1965.

[36] Blevins, R. D.: Applied Fluid Dynamics Handbook. Krieger Publishing, Malabar, FL, 2003.

[37] Embry-Riddle Aeronautical University: Bluff body flows. Introduction to Aerospace Flight Vehicles, Eagle Publishing, Daytona Beach, FL. https://eaglepubs.erau.edu/introductiontoaerospaceflightvehicles/chapter/bluff-body-flows/

[38] Russell, C.; Jung, J.; Willink, G.; and Glasner, B.: Wind Tunnel and Hover Performance Test Results for Multicopter UAS Vehicles. NASA/TM–2018–219758, NASA Ames Research Center, February 2018.

[39] Morris, D.; and Pretty, C.: An experimental study of drag coefficients of a quadrotor airframe. Proceedings of the ASME 2023 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC-CIE), Boston, MA, 2023. doi:10.1115/DETC2023-114020.

[40] Hattenberger, G.; Bronz, M.; and Condomines, J.-P.: Evaluation of drag coefficient for a quadrotor model. Journal of Unmanned Vehicle Systems, 2023. https://enac.hal.science/hal-03924053/file/17568293221148378.pdf

[41] Aerodynamic Efficiency Degradation in Multirotor UAVs. TechRxiv preprint, 2025. doi:10.36227/techrxiv.175018901.11809818/v1.

[42] Rotor Aerodynamic Interaction Effects for Multirotor Unmanned Aircraft Systems in Forward Flight. Thesis, RMIT University, Melbourne, Australia, 2020. https://research-repository.rmit.edu.au/articles/thesis/Rotor_aerodynamic_interaction_effects_for_multirotor_unmanned_aircraft_systems_in_forward_flight/27598599

[43] T-MOTOR: NS62×24 carbon-fibre propeller and U15XXL KV29 propulsion system specifications. https://store.tmotor.com/product/ns62x24-manned-aircraft-carbon-fiber.html

[44] Accurate Predictions of Rotor Hover Performance at Low and High Disc Loadings. University of Glasgow. http://eprints.gla.ac.uk/133234/

[45] Leishman, J. G.; and Syal, M.: Figure of merit definition for coaxial rotors. Journal of the American Helicopter Society, vol. 53, no. 3, July 2008, pp. 290–300. doi:10.4050/JAHS.53.290.

[46] Copernicus Sentinel-2 (processed by ESA): MSI Level-2A BOA Reflectance Product, Collection 1. European Space Agency, 2021. https://sentinels.copernicus.eu/sentinel-data-access/sentinel-products/sentinel-2-data-products/collection-1-level-2a

[47] Zanaga, D.; Van De Kerchove, R.; Daems, D.; De Keersmaecker, W.; Brockmann, C.; Kirches, G.; Wevers, J.; Cartus, O.; Santoro, M.; Fritz, S.; Lesiv, M.; Herold, M.; Tsendbazar, N.-E.; Xu, P.; Ramoino, F.; and Arino, O.: ESA WorldCover 10 m 2021 v200. Zenodo, 28 October 2022. doi:10.5281/zenodo.7254221.

[48] Lang, N.; Jetz, W.; Schindler, K.; and Wegner, J. D.: A high-resolution canopy height model of the Earth. Nature Ecology and Evolution, vol. 7, no. 11, November 2023, pp. 1778–1789. doi:10.1038/s41559-023-02206-6.

[49] U.S. Geological Survey: 3D Elevation Program lidar point cloud. Work units NC_Phase5_2018_A18 and OR_McKenzieRiver_1_2021. https://www.usgs.gov/3d-elevation-program

[50] U.S. Geological Survey: National Hydrography Dataset. https://www.usgs.gov/national-hydrography

[51] OpenStreetMap contributors: OpenStreetMap. Open Database License. https://www.openstreetmap.org/copyright

[52] U.S. Forest Service: Surface Ownership Parcels (BasicOwnership). FSGeodata Clearinghouse. https://data.fs.usda.gov/geodata/

[53] Ronneberger, O.; Fischer, P.; and Brox, T.: U-Net: convolutional networks for biomedical image segmentation. Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015), Lecture Notes in Computer Science vol. 9351, Springer, Cham, 2015, pp. 234–241. doi:10.1007/978-3-319-24574-4_28.

[54] He, K.; Zhang, X.; Ren, S.; and Sun, J.: Deep residual learning for image recognition. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778. doi:10.1109/CVPR.2016.90.

[55] Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P.: Focal loss for dense object detection. IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2980–2988.

[56] Salehi, S. S. M.; Erdogmus, D.; and Gholipour, A.: Tversky loss function for image segmentation using 3D fully convolutional deep networks. Machine Learning in Medical Imaging (MLMI 2017), Lecture Notes in Computer Science vol. 10541, Springer, Cham, 2017, pp. 379–387. doi:10.1007/978-3-319-67389-9_44.

[57] Roberts, D. R.; Bahn, V.; Ciuti, S.; Boyce, M. S.; Elith, J.; Guillera-Arroita, G.; Hauenstein, S.; Lahoz-Monfort, J. J.; Schröder, B.; Thuiller, W.; Warton, D. I.; Wintle, B. A.; Hartig, F.; and Dormann, C. F.: Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. Ecography, vol. 40, no. 8, August 2017, pp. 913–929. doi:10.1111/ecog.02881.

[58] Loshchilov, I.; and Hutter, F.: Decoupled weight decay regularization. International Conference on Learning Representations (ICLR), 2019.

[59] Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q.: On calibration of modern neural networks. Proceedings of the 34th International Conference on Machine Learning (ICML), PMLR vol. 70, 2017, pp. 1321–1330.

[60] Zadrozny, B.; and Elkan, C.: Transforming classifier scores into accurate multiclass probability estimates. Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2002, pp. 694–699.

[61] Lakshminarayanan, B.; Pritzel, A.; and Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.

[62] Sutton, R. S.; and Barto, A. G.: Reinforcement Learning: An Introduction. 2nd ed., MIT Press, Cambridge, MA, 2018.

[63] Towers, M.; Kwiatkowski, A.; Terry, J.; Balis, J. U.; De Cola, G.; Deleu, T.; Goulão, M.; Kallinteris, A.; Krimmel, M.; KG, A.; Perez-Vicente, R.; Pierré, A.; Schulhoff, S.; Tai, J. J.; Tan, H.; and Younis, O. G.: Gymnasium: a standard interface for reinforcement learning environments. arXiv:2407.17032, 2024. https://gymnasium.farama.org/

[64] Mnih, V.; Kavukcuoglu, K.; Silver, D.; et al.: Human-level control through deep reinforcement learning. Nature, vol. 518, no. 7540, 26 February 2015, pp. 529–533. doi:10.1038/nature14236.

[65] van Hasselt, H.; Guez, A.; and Silver, D.: Deep reinforcement learning with double Q-learning. Proceedings of the AAAI Conference on Artificial Intelligence, vol. 30, no. 1, 2016, pp. 2094–2100. doi:10.1609/aaai.v30i1.10295.

[66] Huber, P. J.: Robust estimation of a location parameter. Annals of Mathematical Statistics, vol. 35, no. 1, 1964, pp. 73–101. doi:10.1214/aoms/1177703732.

[67] Hansen, M. C.; Potapov, P. V.; Moore, R.; Hancher, M.; Turubanova, S. A.; Tyukavina, A.; Thau, D.; Stehman, S. V.; Goetz, S. J.; Loveland, T. R.; Kommareddy, A.; Egorov, A.; Chini, L.; Justice, C. O.; and Townshend, J. R. G.: High-resolution global maps of 21st-century forest cover change. Science, vol. 342, no. 6160, 15 November 2013, pp. 850–853. doi:10.1126/science.1244693.

Use of image data permitted with attribution by Esri and its data providers per the Esri Master Agreement, https://assets.esri.com/content/dam/esrisites/en-us/media/legal/ma-full/ma-full.pdf