aerospace

MDPI

Article

Hybrid Strategy Selection and Polygon Decomposition for Energy-Aware 2D UAV Coverage Path Planning

Aramais Tyshchenko [1], Igor Iljin [2], Viktor Masalskyi [2], Oleksandr Masalskyi [3] and Vytautas Bučinskas [2,*]

[1] School of Engineering, University of Warwick, Coventry CV4 7AL, UK; aramais.tyshchenko@warwick.ac.uk
[2] Department of Mechatronics, Robotics and Digital Manufacturing, Faculty of Mechanics, Vilnius Gediminas Technical University, Plytinės g. 25, LT-10105 Vilnius, Lithuania; igor.iljin@vilniustech.lt (I.I.); viktor.masalskyi@vilniustech.lt (V.M.)
[3] Department of Electronic Systems, Vilnius Gediminas Technical University, Plytinės g. 25, LT-10105 Vilnius, Lithuania; oleksandr.masalskyi@vilniustech.lt
* Correspondence: vytautas.bucinskas@vilniustech.lt

Abstract

Energy is the binding constraint on unmanned aerial vehicle survey missions, yet most coverage path planners commit to a single sweep pattern regardless of region shape. Our central contribution is to select decomposition granularity and the per-cell coverage pattern by a single physics-based energy criterion. A boustrophedon cellular decomposition with a controlled-concavity merge sets granularity; for each sub-polygon, the planner evaluates a minimum-width lawnmower and a shape-mimicking spiral and keeps the cheaper, with the undecomposed single-sweep solution always in the search space. On 50 benchmark polygons, at a 90% coverage floor, the hybrid reduces energy over a spiral baseline by a median 4.8% (ratio of means 11%) and is never dearer than either fixed pattern at floors of 90–99%. Which pattern is cheaper depends on geometry—lawnmower and spiral are not separable in aggregate (Wilcoxon $p = 0.47$). A boundary-reflective zigzag for double coverage achieves 46% fewer turns and 18% less energy than boustrophedon 2×. The hybrid keeps the lowest mean energy under $\pm$20% parameter perturbation and on larger fields with wider row spacing.

Keywords: coverage path planning; unmanned aerial vehicle; energy optimisation; polygon decomposition; strategy selection; boustrophedon; spiral coverage; benchmark evaluation

Check for updates

Academic Editors: Sergio Esteban, Guilherme V. Raffo, Daniel Neri Cardoso and Marcelo Alves Dos Santos

Received: 8 September 2026
Revised: 9 October 2026
Accepted: 19 October 2026
Published: 22 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Coverage path planning (CPP) is the problem of determining a path for a UAV such that every point in a given region is observed by the vehicle's sensor. In UAV applications such as precision agriculture, search and rescue, environmental monitoring, and infrastructure inspection, efficient coverage directly translates to reduced mission time, energy consumption, and operational cost [1].

The fundamental challenge in 2D CPP lies in balancing multiple competing objectives: complete sensor coverage must be achieved while minimising path length (and therefore energy consumption), the number of directional changes (turns), and total mission time. These objectives often conflict: minimising path length may increase turns, while reducing turns may leave coverage gaps. Concretely, we judge a coverage plan by two criteria. Completeness is the fraction of the region actually swept by the sensor footprint—a good plan leaves no gaps, whereas a poor one leaves uncovered slivers along boundaries or inside narrow features (quantified in Section 3). Cost is the energy required to fly the plan; on a multirotor, much of it is spent in the decelerate–rotate–accelerate maneuver at each turn, so

a plan that covers the same area with fewer turns is preferred. The aim is the lowest-energy plan that attains the required completeness at a given overlap factor.

Formally, given a simple polygon $P \subset \mathbb{R}^2$ (possibly non-convex), a sensor footprint width $w > 0$ (a ground length), an overlap factor $\rho \in [0, 1)$, and a required coverage completeness $C_{\min}$, the goal is to find the minimum-energy continuous path π^* that attains the required completeness:

$$\pi^* = \arg\min_{\pi} E_{\text{total}}(\pi) \quad \text{subject to} \quad C(\pi, w) \geq C_{\min} \tag{1}$$

where $C(\pi, w)$ denotes the fraction of P swept by the sensor footprint along π. We use $C_{\min} = 90\%$ as the operating-point floor on the benchmark suite (Section 2.5) and repeat the comparison at 95%, 98%, and 99%. Comparisons are made at each strategy's own lowest-energy configuration meeting this floor, and no strategy is credited with a cheap-but-undercovering solution.

Which plan is cheapest depends on the regime, which two ratios set. The 50 benchmark polygons have $A/w^2 = 120$–714 and 4–36 sweep rows at $\rho = 0$. A multirotor needs $v_{\text{cruise}}^2/a_{\max} = 40$ m to reach the 10 m/s cruise speed at 2.5 m/s^2 and stop again, whereas in the base suite, the rows are 7–49 m long (median 15 m), so on 49 of the 50 polygons, no run between turns reaches cruise speed, and the turn/maneuver share of the energy is about 93%. Section 3.8 repeats the comparison with every polygon scaled to row spacings of 5 and 10 m (fields of 0.3–7.1 ha) at cruise speeds of 5 and 10 m/s, where that share is 30–75%, so the evaluation spans both turn-dominated and cruise-significant regimes.

Existing CPP methods predominantly employ a single coverage strategy, typically boustrophedon (lawnmower) sweeps [2], applied uniformly across all polygon cells after decomposition. Which of the two patterns is cheaper is contested: Cabreira et al. [3] found an energy-aware spiral cheaper than an energy-aware back-and-forth pattern on every one of more than 3500 simulated polygons, whereas against a turn-minimising minimum-width lawnmower, we find each pattern cheaper on some shapes (Section 3.2). This motivates per-polygon strategy selection based on local geometry and energy cost. The gap this leaves is a planner in which decomposition granularity and per-cell coverage pattern are both selected by the same energy criterion: prior methods fix one and tune the other. This paper's main contribution is exactly that coupling—a single energy-minimising formulation in which how finely a region is decomposed and which pattern sweeps each cell are both selected by one energy search, with the undecomposed single sweep always available as a candidate, evaluated on 50 polygons (the feature-level comparison against representative prior methods and the quantitative comparison with three planners are given in Section 4.1).

The individual ingredients we build on are established: energy-aware back-and-forth and spiral planning [3,4], concavity-limited decomposition [5,6], and physics-based multirotor power models [4,7]. Our contribution is to combine and re-frame them, specifically:

1. A unified energy-minimising formulation in which decomposition granularity (via a controlled-concavity angle threshold θ applied to a boustrophedon cellular decomposition) and per-cell pattern (lawnmower or spiral) are outcomes of a single energy criterion rather than hand-set choices. Each cell is swept by a minimum-width boustrophedon, with the undecomposed single-sweep solution guaranteed to lie in the search space (Sections 2.1 and 2.4). Prior constriction and approximate-convex decompositions fix granularity before planning [5,6,8], and prior energy-aware planners optimise one pattern over a given region or decomposition [3,4,9]. We instead search granularity under the same energy criterion that makes the per-cell lawnmower/spiral choice.
2. A boundary-reflective diagonal zigzag for double coverage on convex and mildly concave cells is presented, which lowers turns and energy relative to a boustrophedon 2× baseline (Section 3.10).

3. A systematic evaluation is conducted across 50 benchmark polygons, more than the 1 to 33 instances reported by the decomposition-based CPP methods it is compared against (Section 4.1), at four coverage floors, with a one-at-a-time ablation of the method's components, a comparison with three other planners, and checks of its sensitivity to the energy-model parameters and to the form of the turn-energy model (Sections 3, 3.7 and 4).

Searching a decomposition under a mission cost is itself established [10], as are turn-reducing decompositions [8,11,12], but these sweep every cell with one back-and-forth pattern. What is new here is narrower: the per-cell lawnmower/spiral choice is made inside the energy search over granularity and slab orientation, with the undecomposed sweep always a candidate; the merge applies the largest turn gain first under an interior-angle bound, so it can form mildly concave cells; and the zigzag applies boundary reflection to implicit double coverage. The transit-inclusive minimum of Section 2.4 is a safeguard, and the grid search and 50-polygon suite are an evaluation protocol.

Across a geometrically diverse benchmark, the study is organised around three questions: Which coverage pattern minimises energy for a given polygon? How much does selecting decomposition granularity and pattern under one energy criterion save over a fixed single-strategy baseline? Finally, on which shapes does the choice actually matter? Our results answer that on the base suite at the 90% floor, lawnmower has the lower mean energy overall and is the lower-energy whole-polygon pattern on most shapes (28 of 50). The contour-following spiral wins the other 22. The energy-minimising plan usually involves little or no decomposition, and the resulting hybrid reduces energy by a median 4.8% over a spiral baseline while never exceeding the lawnmower baseline at that floor.

Section 2 describes the controlled-concavity decomposition, the lawnmower and spiral path generators, the physics-based energy model, the strategy-selection mechanism, and the benchmark suite and grid-search protocol. Section 3 reports the strategy comparison, the comparisons between strategies at four coverage floors, parameter effects, an ablation of the method's components, category effects, coverage analysis, sensitivity analysis, scalability, and the double-coverage study. Section 4 positions the work against the literature, and Section 5 concludes and states the study's limitations.

1.1. Related Work

1.1.1. Coverage Path Planning Foundations

Boustrophedon cellular decomposition, introduced by Choset [2], divides free space into cells using critical points and covers each cell with back-and-forth motions. The most comprehensive CPP survey, by Galceran and Carreras [1], identifies energy awareness and adaptive decomposition as key research gaps. Huang [11] established the optimal line-sweep approach for convex polygons, and Bochkarev and Smith [13] refined the selection of sweep directions across a decomposition to minimise total turns; our per-cell minimum-width sweep direction, chosen by rotating calipers as the direction that minimises the number of sweep lines, follows this line of work. Early evidence for shape-dependent strategy selection came from Oksanen and Visala [14], who showed that polygon shape strongly influences algorithm selection.

1.1.2. Decomposition-Based Approaches

The Hertel–Mehlhorn algorithm [15] removes inessential diagonals from a triangulation, giving a convex partition with at most four times the minimum number of pieces. Relaxing strict convexity (approximate convex decomposition) yields considerably fewer pieces than an exact convex decomposition [6], and Brown and Waslander [5] explicitly limit cell count by decomposing only at concavity "constrictions" rather than at every reflex vertex. This is the prior work closest in spirit to our concavity-tolerant merging, although

they fix the decomposition rather than treating its granularity as an energy-optimisation variable. A concavity-eliminating convex decomposition with flight-direction selection and sub-region merging has likewise been used to reduce sub-region count and turns for back-and-forth coverage [8]. There, the decomposition is a turn-minimising pre-processing step rather than a granularity knob tuned by the coverage cost itself. Other work has compared decomposition strategies for UAV surveys and again found strategy dependence on geometry [16], proposed convex decomposition that extends interior edges [17], and introduced the R-DFS width-surrogate approach [18]. Li et al. [12] reduced the coverage problem to a polygon-width computation with an $O(n)$ width algorithm, decomposed the region to minimise the sum of subregion widths, and connected subregions through a joint-point scheme so as to minimise the number of turns. We reimplement its decomposition as a baseline, chaining the cells with our own tour rather than its joint-point scheme (Section 4). Conversely, Kusnur and Likhachev [19] argue that decomposition can itself be suboptimal and pursue a decomposition-free formulation. Our search side-steps this by retaining the undecomposed sweep as a candidate.

1.1.3. Energy-Aware Path Planning

Di Franco and Buttazzo [4] derived a measurement-based multirotor energy model covering acceleration, deceleration, turning, and hovering and used it to plan energy-aware back-and-forth paths. For fixed-wing surveys in steady wind, choosing the decomposition and sweep directions under a wind-aware flight-time cost shortened flights by roughly 5–25% in the showcase fields of [10], and the flight-time-optimal sweep angle has no simple relation to the wind direction [20], while an MILP formulation of energy-aware CPP guarantees optimality but scales poorly [21]. Hong et al. [22] developed a flight-validated power model for multirotors, and Datsko et al. [9] planned multi-UAV boustrophedon coverage at the minimum-energy flight speed of a first-principles power model validated in flight, choosing the decomposition orientation by a sum-of-subregion-heights cost. We instead fix the cruise speed and treat decomposition granularity and the per-cell coverage pattern as the energy-optimisation variables. Further approaches include an integrated system-dynamics treatment [23], deep reinforcement learning for energy-constrained CPP [24], and adaptive sweep spacing coupled with a power-estimation module [25]; a survey confirms that energy-aware decomposition-based methods remain most practical for single-UAV missions [26]. Closest to the present work in judging coverage patterns by energy, Cabreira et al. [3] gave a spiral optimal speeds on its straight segments and found it cheaper than an energy-aware back-and-forth pattern on every one of more than 3500 simulated polygons, while Jensen-Nau et al. [27] generate energy-constrained partial-coverage paths by a Voronoi-based method that are cheaper than a lawnmower path; each optimises one pattern rather than choosing between patterns. The inward spiral generated by iterative polygon offsetting that we use is itself a known construction [28] and our contribution is not the spiral, but the selection mechanism around it. Khaneghaei et al. [29] refine an offline back-and-forth sweep online by model predictive control with Dubins smoothing, so that a fixed-wing robot can track it among obstacles; energy enters their controller's cost but not the choice of pattern or decomposition, so such an execution layer is complementary to the planner presented here. A multirotor, unlike a fixed-wing vehicle, can decelerate, yaw in place, and re-accelerate, so we treat its dynamics as a cost rather than a constraint (Section 2.3).

The aim of this work is to design and demonstrate an energy-minimising coverage planner in which the granularity of the polygon decomposition and the coverage pattern flown inside each resulting cell are selected by a single physics-based energy criterion, rather than being fixed by hand before any coverage cost is known. In general, this work offers several contributions to decomposition-based CPP and to the survey tasks that rely on

it, such as precision agriculture, infrastructure inspection, and environmental monitoring, in which endurance, not sensor capability, is the binding constraint. Most existing planners commit to one sweep pattern, most often a boustrophedon, and apply it uniformly over a decomposition whose granularity is often set before the coverage cost is evaluated; that ordering makes the two decisions dominating multirotor energy consumption mutually blind, and it imposes a decomposition even on regions that a single sweep covers more cheaply. The proposed method addresses two coupled problems: the per-sub-polygon choice between a minimum-width lawnmower and a contour-following spiral, and the choice of whether and how finely a non-convex region should be split before it is swept, with the undecomposed single sweep always retained in the search space. Our approach differs from the closest energy-aware coverage work [3,27] in two respects: the lawnmower-versus-spiral choice is made per sub-polygon after a tunable-concavity decomposition rather than fixed in advance; and its advantage over the lawnmower baseline is reported as a dominance count rather than as a stochastic comparison. The planner is verified over a 7200-configuration grid search on 50 benchmark polygons spanning five geometric categories, compared on all 50 polygons with three other planners [12,30], and stress-tested by ±20% perturbation of five energy-model parameters and by two alternative turn-cost models.

2. Materials and Methods

We pose the coverage task of Equation (1)—find the minimum-energy path that meets a required coverage completeness at a given overlap factor—and solve it by simulation: for each benchmark polygon, we generate candidate paths under a physics-based energy model and select the configuration that minimises energy, and then report how the comparisons respond to the model's parameters and to the form of its turn-cost term. The proposed approach consists of four components: (1) controlled-concavity polygon decomposition, (2) path generation via lawnmower or spiral algorithms, (3) physics-based energy evaluation, and (4) strategy selection. The remainder of this section details each component and then the benchmark suite, grid-search protocol, and metrics used to evaluate them. All planning, energy evaluation, and analysis were implemented in Python 3.11 (Shapely 2.1 for geometry, NumPy 2.3 and SciPy 1.16 for numerics and statistics).

2.1. Controlled-Concavity Decomposition

Given a possibly non-convex polygon P, the decomposition proceeds in two stages. First, a sweep line is passed across P, and the polygon is sliced at every vertex event, a trapezoidal refinement of Choset's boustrophedon cellular decomposition [2], which cuts only at critical points; the merge below removes the redundant cuts. Between two consecutive events, the boundary has no vertices, so each cross-section is bounded above and below by a single straight edge and the base cells' only internal cut-edges are aligned with the sweep. This matters for the subsequent sweep: a triangulation-based decomposition introduces diagonal internal cut-edges, and a per-cell boustrophedon that orients to such a cell then produces spurious diagonal passes.

Second, a pair of adjacent base cells is eligible to merge if their union is a simple polygon whose interior angles all remain at or below the threshold θ—an almost-convex criterion in the spirit of Hertel–Mehlhorn's convex merging [15]:

$$P_i \cup P_j \text{ is a simple polygon} \quad \text{and} \quad \max_k \alpha_k(P_i \cup P_j) \leq \theta \tag{2}$$

Among the eligible pairs, we merge greedily by estimated turn reduction rather than accepting the first feasible merge. A boustrophedon over a cell of minimum width $w_{\min}$ needs $r = \lceil w_{\min}/w \rceil$ sweep rows and $\approx 2(r-1)$ U-turns, and two separate cells additionally cost

≈2 turns for the transit between them, so merging P_i, P_j into $P_i \cup P_j$ changes the turn count by $g = 2\,(r_i + r_j - r_{ij})$ (the transit term is a structural count, not a tuned constant). At each step, the eligible merge with the largest positive g is applied, until no positive-gain merge remains. This drives the decomposition toward the fewest turns—matching the low cell count of a minimum-width-sum convex decomposition [12]—while the θ gate still permits merges into mildly concave cells that a strictly convex scheme cannot form. The base decomposition is exhaustive but over-fragments the region; it is this bounded-concavity, turn-aware merge—not the base decomposition—that is the tunable contribution.

The decomposition is additionally searched over its orientation. Boustrophedon slabbing along a single global direction leaves axis-aligned cells; we instead try the slab direction at 5° steps (0–175°) and keep the orientation whose resulting decomposition minimises the chained energy of the per-cell-mixed plan at the coverage floor. This lets the cells—and hence their per-cell minimum-width sweeps—align to an inclined region axis, the same edge-parallel principle behind Li et al.'s enhanced ECD [12], and is what lets the method produce inclined passes on concave shapes rather than only axis-aligned ones. The search is run in a coverage-preserving mode: an inclined orientation is adopted only when the coverage of the per-cell-mixed plan stays within 1.5 percentage points of its axis-aligned value and above the floor, so an inclined decomposition can trade at most 1.5 points of coverage for its energy saving; where no orientation meets this rule, the orientation whose mixed plan covers most is kept. Orientation is shared across the three strategies at each configuration, and convex regions and the undecomposed $\theta = 360°$ case have no orientation degree of freedom.

Finally, any cell too thin to sweep—one whose minimum width is below the sensor footprint, so the half-footprint inset would remove its only row—is absorbed into the adjacent cell with the longest shared boundary. This leaves every retained sub-polygon wide enough for at least one sweep row and removes the redundant fragments an inclined slabbing can shed at oblique corners, reducing the cell count without loss.

The threshold θ controls decomposition granularity: $\theta = 180°$ permits only convex merges; $\theta = 200$–$260°$ tolerates mild concavity, producing fewer sub-polygons; and at $\theta = 360°$, the merge is bypassed—the intact polygon P enters the search space directly as a single-cell candidate. Figure 1 illustrates this effect on a representative polygon. Equation (2) would admit every adjacent pair at $\theta = 360°$, but a merge also needs a positive estimated gain g, so the greedy merge does not reliably reassemble P (on the L-shape, for example, axis-aligned slabs stay split while slabs cut at 135° merge into one cell); injecting P directly keeps the undecomposed sweep in the search space on concave regions too.

Once the decomposition has fixed the set of sub-polygons, each cell is covered by one of the two path generators described next, whose per-cell cost the strategy selection of Section 2.4 then compares.

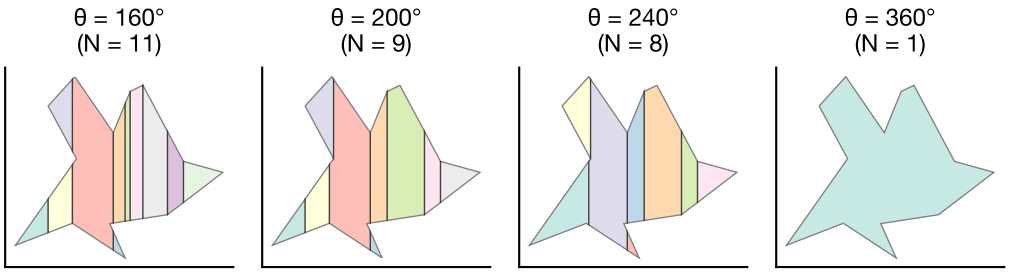


Figure 1. Effect of angle threshold θ on polygon decomposition. Increasing θ reduces sub-polygon count N at the cost of allowing mild concavity. Each random colour marks one sub-polygon.

2.2. Path Generation Algorithms

2.2.1. Lawnmower (Boustrophedon) Coverage

The lawnmower algorithm generates back-and-forth sweeps parallel to the cell's minimum-width support line, the direction that minimises the number of sweep lines—and hence U-turns—for a given footprint [11,12]: (1) the minimum-width direction is found by rotating calipers on the cell's convex hull, and the cell is rotated so that direction is horizontal and inset by $\frac{w}{2}$ on all sides; (2) sweep lines are generated at spacing $\Delta y = w(1-\rho)$, the first and last centre-lines placed $w/2$ from the boundary so the footprint just reaches the edges parallel to the sweep; (3) the swept segments are connected in alternating directions, and where a sweep line is split by a concavity, the disjoint strands are joined by a connector routed through the cell interior, so the trajectory stays inside the polygon and the connecting turns are counted. The minimum-width orientation is the textbook turn-minimising choice for a convex cell and, applied to the trapezoidal cells above, gives clean axis-aligned sweeps free of the diagonal artefacts a longest-edge rule would inherit from internal cut-edges.

2.2.2. Shape-Mimicking Spiral Coverage

The spiral algorithm iteratively shrinks the polygon by $\Delta = w(1-\rho)$ using a negative buffer (Minkowski difference), forming concentric loops. Two ways of joining successive loops are considered: a continuous inward spiral (the loops spliced into one smooth path, few turns) and concentric rings joined by short aligned transitions (more turns, and no splice that could cut a chord across a concave cell). To avoid handicapping the spiral, the spiral baseline uses the cheaper of the two whenever the continuous spliced path is valid—i.e., it stays inside the cell (within a small tolerance) and covers essentially as well as the rings (within two points); otherwise, it falls back to rings. This is the same coverage-gated spiral the hybrid can select per cell, so the hybrid is compared against the best spiral the method can produce, not a weaker ring-only baseline. The rule is applied identically wherever a spiral is costed, so the reported energy matches the flown path. Spiral paths produce fewer turns on compact polygons but more on elongated or angular geometries.

2.3. Physics-Based Energy Model

Total mission energy is the sum of cruise, maneuver (the acceleration/deceleration lift and kinetic work), per-turn (yaw or banked), and systems components:

$$E_{\text{total}} = E_{\text{cruise}} + E_{\text{maneuver}} + \sum_{\text{turns}} E_{\text{turn}} + E_{\text{systems}} \tag{3}$$

Cruise power follows the standard reduced form of the rotary-wing forward-flight power curve [4,31], with a constant term (c_1), a speed-squared blade-profile term ($c_3 v^2$), and the high-speed asymptote of the induced-power term (c_4/v, decreasing with airspeed):

$$P(v) = c_1 + c_3 v^2 + \frac{c_4}{v} \tag{4}$$

This produces the characteristic U-shaped power curve with a finite minimum-power airspeed $v^* = (c_4/2c_3)^{1/3}$ (Figure 2a). Two simplifications are made: the $c_3 v^2$ term is the blade-profile forward-flight correction (not vehicle parasite drag), and the cubic parasite term ($\propto v^3$) is omitted, which is acceptable at the modest speed used here; and because the c_4/v-induced term is a forward-flight asymptote that diverges as $v \to 0$, hover is defined off-curve (the separate P_{hover} below) rather than read from Equation (4). This constant-plus-v^2-plus-$1/v$ form is the parasite-dropped reduction of the rotary-wing propulsion-power model of Zeng et al. [31]; fuller small-multirotor models that additionally resolve body pitch and rotor drag reduce to the same functional family once the small cubic parasite term is dropped, so retaining only these three terms is a reasonable approximation at the low cruise speed used here—and in the base suite, the strategy comparison is dominated

by maneuver energy rather than cruise (Section 3), so its relative rankings are insensitive to the exact cruise-curve shape; Section 3.8 repeats the comparison where cruise takes 19–62% of the energy. The coefficients (Table 1) are set within the family of flight-validated small-multirotor power models [4,7] for a ∼1.5 kg Phantom-4-class quadrotor, giving a hover draw of ∼140 W (momentum-theory and battery-endurance estimates for this class span ≈140–180 W; we take the lower, optimistic end), $v^* \approx 4.6$ m/s, and $P(v^*) \approx 97$ W (≈0.69 of hover). The survey cruise speed (10 m/s) lies above v^*, a standard operational choice for shorter mission time and higher area rate. Within each comparison, the cruise speed is the same for all strategies; Section 3.8 repeats the comparison at 5 m/s, where the margins and the order of the two fixed patterns change.

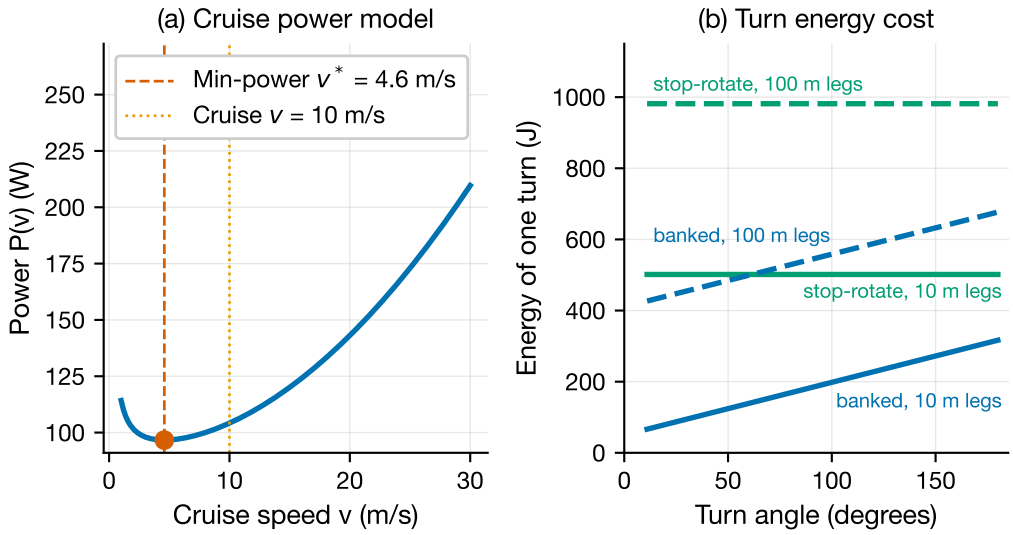


Figure 2. Energy model characteristics. (**a**) Cruise power vs. airspeed; the dot marks the minimum-power speed v^* and its power $P(v^*)$. (**b**) Energy added by one turn between two legs of 10 m or 100 m, under the stop–rotate model and the banked-turn ablation, computed from the model.

Table 1. Energy model parameters.

Parameter	Symbol	Value	Unit
Vehicle mass	m	1.5	kg
Motor efficiency	η	0.85	–
Constant power term	c_1	88.2	W
Blade-profile coefficient	c_3	0.134	W·s^2/m^2
Induced power	c_4	25.9	W·m/s
Hover power	P_{hover}	140	W
Systems power	P_{sys}	10	W
Yaw time per turn	t_{yaw}	1.0	s
Max acceleration/deceleration	$a_{\max}$	2.5	m/s^2
Max bank angle (ablation)	$\phi_{\max}$	45	degrees
Cruise velocity	v_{cruise}	10	m/s
Turn velocity (ablation)	v_{turn}	4	m/s

The maneuver and turn energy is computed by flying the vehicle along the path with a segment-length-aware trapezoidal speed profile. Between successive stops—the path endpoints and every significant heading change (turn angle $> 10^\circ$ between consecutive path segments, including the turn where an inter-cell transit joins the first row of the next cell)—it accelerates from the entry speed toward v_{cruise}, cruises only if the run is long enough, and then decelerates to the exit speed, bounded by the acceleration/deceleration

limit $a_{\max}$ (Table 1); the peak speed v_{peak} is capped by the run length, so on short rows, the vehicle never reaches cruise. Over a run of length L, the energy is

$$E_{\text{run}} = P(v_{\text{cruise}})\, t_{\text{cruise}} \;+\; P_{\text{hover}}\,(t_{\text{acc}} + t_{\text{dec}}) \;+\; \tfrac{1}{2}\tfrac{m}{\eta}\,(v_{\text{peak}}^2 - v_{\text{in}}^2), \tag{5}$$

i.e., cruise power over any constant-speed phase, weight-supporting lift power ($\approx P_{\text{hover}}$) over the low-speed acceleration and deceleration phases, and the kinetic motor work to reach the run's peak speed from its entry speed v_{in} (zero after a stop). Charging lift power throughout the maneuver—not only during the heading change itself—follows the "hover power below the optimal speed" treatment of Datsko et al. [9] and the phase-integrated power accounting of Di Franco and Buttazzo [4]. For the primary multirotor model, each significant heading change is a stop–rotate: the vehicle stops and yaws in place at hover power for a nominal time t_{yaw} (adding $P_{\text{hover}} t_{\text{yaw}}$ per turn). As a structural robustness check, we also evaluate a fixed-wing-style coordinated banked turn at maximum bank angle $\phi_{\max}$:

$$E_{\text{turn}}^{\text{banked}} = \frac{P_{\text{hover}}}{\cos\phi_{\max}} \cdot \frac{R\delta}{v_{\text{turn}}}, \tag{6}$$

where $R = v_{\text{turn}}^2 / g \tan\phi_{\max}$ is the turn radius and δ the turn angle, in which the vehicle only slows to v_{turn}, the slow-down and re-acceleration being charged by Equation (5) with entry and exit speed v_{turn}; both models are evaluated on identical paths, and Section 3.7 reports how the comparisons change between them. Because the banked arc's 1.63 m radius does not fit the sub-metre corridors of the base suite, this ablation is an optimistic lower bound on turn energy rather than a feasible model of a reversal. A direct consequence of this accounting is that at the region scales of the base suite, the flight is maneuver-dominated: no run between turns reaches cruise speed on 49 of the 50 polygons (Section 1), so the decelerate–yaw–accelerate maneuver accounts for ≈93% of mission energy and cruise for 0.05% of the pooled energy at the operating points. This share falls as rows lengthen, which we quantify with a scale sweep and at row spacings of 5–10 m (Section 3.8).

At the row spacings of the base suite, a reversal is necessarily a near-stop. A heading reversal between rows a distance d apart can be flown inside the corridor only on a turn of radius at most $d/2$, which limits the speed through it to $\sqrt{a_{\text{lat}}\, d/2}$, i.e., 0.66–1.12 m/s for $a_{\text{lat}} = a_{\max}$ and $d = 0.35$–1.0 m (1.3–2.2 m/s even at a 1 g lateral acceleration), against a 10 m/s cruise; a wide turn at cruise speed would need a radius of $v_{\text{cruise}}^2 / a_{\text{lat}} = 40$ m and would leave the polygon, which the planner does not allow. The model charges this near-stop as a full stop, as in rest-to-rest coverage planning flown on real multirotors [4,32,33]. At the 5–10 m spacings of Section 3.8, the same bound allows 2.5–3.5 m/s through a reversal, and about 5–7 m/s at the 1 g lateral acceleration of the banked-turn model, so the full stop there overstates the turn cost, and that section also reports the banked-turn model. The in-place yaw added at each stop is a modelling assumption rather than a requirement—a vehicle with a gimbal-stabilised sensor can translate without yawing—and Section 3.7 also reports the comparison without it. Conversely, the blanket 10° threshold over-charges shallow bends between long legs, where a vehicle could carry speed; such bends occur mostly in spiral paths, which is why the spiral comparison is also reported under the banked-turn model.

Systems energy captures constant avionics draw, $E_{\text{sys}} = P_{\text{sys}} \cdot t_{\text{mission}}$, integrated over the full mission time (including all maneuver phases, so the time and energy ledgers describe the same flight). The launch and landing accelerations are simply the first and last runs of the trapezoidal profile above (entry/exit speed zero), so no separate ramp term is required.

Table 1 lists all model parameters, representative of a mid-size survey quadrotor (DJI Phantom 4 RTK class).

These parameters produce the two characteristic curves of the model, plotted in Figure 2: the U-shaped cruise-power curve and the per-turn maneuver cost.

With the decomposition and the two path generators defined and their flight energy quantified, we can now state how the planner chooses among them.

2.4. Strategy Selection

For each sub-polygon, the candidate paths are generated, and the hybrid strategy keeps the cheapest, providing a structurally guaranteed non-negative improvement over either fixed strategy. The candidate set per cell is the minimum-width lawnmower and two spiral variants—the concentric-ring spiral and, when it stays inside the cell, the continuous inward spiral—so $E_{\text{hybrid}} = \min$ over these; on a strictly convex cell, the continuous spiral is always valid, whereas on a mildly concave cell (permitted by the controlled-concavity merge), it is admitted only when the inside-cell test passes, giving the hybrid an option a strictly-convex decomposition cannot offer. To avoid trading coverage for turns, the continuous spiral is admitted per cell only when its coverage is within about two points of the concentric-ring spiral's; otherwise, the hybrid falls back to the ring spiral or lawnmower on that cell. The low-turn continuous spiral is thus kept wherever it is coverage-equivalent, and the more robust ring spiral is used exactly on the cells where a spliced spiral would leave a gap. Once each cell's pattern is fixed, the cells are sequenced by a multi-start greedy nearest-neighbour tour (each tour starting from a different cell, every later cell flown in whichever direction starts nearer), and the lowest-energy of these orderings is kept; this reduces, but does not minimise, inter-cell transit (the exact order is cheaper on 8 of the 50 operating points, where the tour costs at most 3.7% more, Section 3.4). Because the per-cell argmin is taken before chaining, it ignores the selection-dependent inter-cell transit, and a per-cell-cheaper spiral can occasionally lengthen a transit enough to make the chained mixed path cost more than a uniform strategy. To keep the dominance guarantee exact under our chaining, the reported hybrid is therefore the lowest-energy chained candidate among the per-cell-mixed path, the uniform minimum-width lawnmower, and the uniform spiral: the per-cell mixing benefit is retained where it helps, and the transit downside is bounded, so $E_{\text{hybrid}} \leq \min(E_{\text{lawnmower}}, E_{\text{spiral}})$ holds by construction at every configuration. When the operating point is selected, the coverage floor is applied to each of the three candidates before the cheapest is taken, so a candidate that meets the floor is never hidden behind a cheaper one that does not. The uniform candidates are the fixed patterns' own plans at the same configuration, so the hybrid's operating point is never dearer than a fixed pattern's operating point that meets the floor (Section 3.2).

Crucially, the decomposition itself is part of the search rather than a fixed pre-processing step. The angle threshold θ is selected from the candidate set (Section 2.5) by the same energy criterion: for each θ, the polygon is decomposed, the per-sub-polygon minimum-energy strategy is chosen, and the configuration with the lowest total energy is retained. Because the candidate set includes $\theta = 360^\circ$, at which the intact polygon is injected directly as a single-cell candidate (Section 2.1), an undecomposed single-sweep solution is always within the search space.

The per-polygon operating points of Section 3 also take the lowest-energy result over the overlap factors of the grid. The threshold and orientation searches are exhaustive over their finite sets, and the per-cell choice is exact over its candidates; the merge (greedy, largest gain first), the transit-free per-cell choice, and the tour are heuristic, and of these, only the tour's gap is measured (Section 3.4). Throughout, "energy-optimal" and "energy optimum" denote the lowest energy among the configurations evaluated, not a global optimum over decompositions, assignments, and visit orders. Because both undecomposed sweeps are among the candidates, no operating point is dearer than an undecomposed sweep that meets the floor: on the L-shape, a single contour-following pass is cheapest and the search returns it, whereas on the complex shape, it returns a multi-cell solution (Section 4). This makes decomposition granularity and per-cell algorithm choice outcomes of a single energy search, not hand-set parameters; within each cell, the sweep follows the minimum-width boustrophedon orientation. Algorithm 1 summarises the full procedure.

Algorithm 1 Energy-minimising decomposition and per-cell strategy selection

Require: polygon P, footprint w, overlap ρ, threshold set Θ, coverage floor $C_{\min}$
Ensure: lowest-energy plan π^* meeting the floor

best $\leftarrow \emptyset$
for all $\theta \in \Theta$ **do** ▷ granularity is searched, not fixed
  if $\theta = 360^\circ$ **then** ▷ bypass the merge: inject intact P (Section 2.1)
    $\{c_i\} \leftarrow \{P\}$ ▷ no slabbing, hence no orientation search
  else
    $\phi^* \leftarrow \arg\min_\phi E\big(\pi_{\text{mix}}(\textsc{AbsorbThinCells}(\textsc{Decompose}(P, \theta, \phi), w))\big)$ over $\phi \in \{0^\circ, 5^\circ, \ldots, 175^\circ\}$, coverage-preserving ▷ per-cell-mixed plan
    $\{c_i\} \leftarrow \textsc{AbsorbThinCells}\big(\textsc{Decompose}(P, \theta, \phi^*), w\big)$ ▷ cells of the chosen orientation
  end if
  for all cells c_i **do**
    $p_i^{\text{lm}} \leftarrow \textsc{MinWidthLawnmower}(c_i, w, \rho)$; $p_i^{\text{sp}} \leftarrow \textsc{Spiral}(c_i, w, \rho)$
    $p_i^{\text{mix}} \leftarrow \arg\min_{p \in \{p_i^{\text{lm}}, p_i^{\text{sp}}\}} E(p)$ ▷ per-cell argmin (transit-free)
  end for
  for all $s \in \{\text{mix}, \text{lm}, \text{sp}\}$ **do**
    $\pi_s \leftarrow \textsc{BestChain}\big(\{p_i^s\}, P\big)$ ▷ multi-start NN tour, in-polygon transits
  end for
  for all $s \in \{\text{mix}, \text{lm}, \text{sp}\}$ **do** ▷ floor applied to each candidate
    if $C(\pi_s) \geq C_{\min}$ **and** $E(\pi_s) < E(best)$ **then**
      $best \leftarrow \pi_s$
    end if
  end for
end for
if $best = \emptyset$ **then** ▷ best effort
  $best \leftarrow$ the highest-coverage π_s over Θ and s
end if
return *best*

2.5. Simulation Setup

This study is a full-factorial simulation: every benchmark polygon is solved under every combination of the controllable settings, and the energy of each resulting path is evaluated with the model of Section 2.3. We sweep three axes—the decomposition angle threshold θ from 160° to 360° (eight levels), the overlap factor ρ from 0.0 to 0.5 (six levels), and the coverage strategy (lawnmower, spiral, hybrid)—over 50 polygons whose areas span 40–260 m^2 and footprint widths 0.35–1.0 m. Two orientation choices must be distinguished: within each cell, the sweep orientation is fixed to that cell's minimum-width direction (no per-cell sweep search), whereas the decomposition slab orientation is searched at 5° steps (Section 2.1), shared across the three strategies. This slab-orientation search is an inner optimisation, run once for each of the $50 \times 8 \times 6 = 2400$ (polygon, θ, ρ) triples and shared by the three strategies, not a fourth grid axis. For each polygon, we then read off the energy-optimal configuration subject to a coverage floor: the lowest-energy setting whose coverage reaches at least 90%, or, where no setting does, the highest-coverage setting (best effort). This selection keeps a strategy from being credited with an operating point that is cheap only because it covers less than the floor; completeness can still differ between strategies above the floor, so coverage is reported beside energy. These 50 per-polygon operating points—the configurations a planner would deploy—are used for all per-polygon comparisons. The 90% value is a screening constraint rather than a coverage target: it is set below the coverage most operating points reach (mean 96.5% for the hybrid, 96.7% for the spiral, and 97.4% for the lawnmower), although six hybrid operating points lie within two points of it. The comparison is repeated at floors of 95%, 98%, and 99% (Section 3.2), with the whole grid re-run at each floor, because the orientation search and the best-effort fallback depend on it. The benchmark suite, the grid, and the metrics are detailed next.

2.5.1. Benchmark Suite

A structured benchmark suite of 50 polygons spans five geometric categories (Table 2). Polygons cover areas from 40 to 260 m^2, with footprint widths from 0.35 to 1.0 m; all 50 are drawn in Appendix A, and Section 3.8 also evaluates them scaled to footprint widths of 5 and 10 m. The multi-concavity "Complex Shape" instance (Figure 3 and Appendix A) reproduces the exact 13-vertex enhanced exact-cellular-decomposition benchmark of Li et al. [12] (their Figure 14), rescaled to the suite's area range, a shape that recurs as a test case across the decomposition-based CPP literature; the remaining polygons are synthetic.

Each of these polygons is solved under every setting of the grid described next.

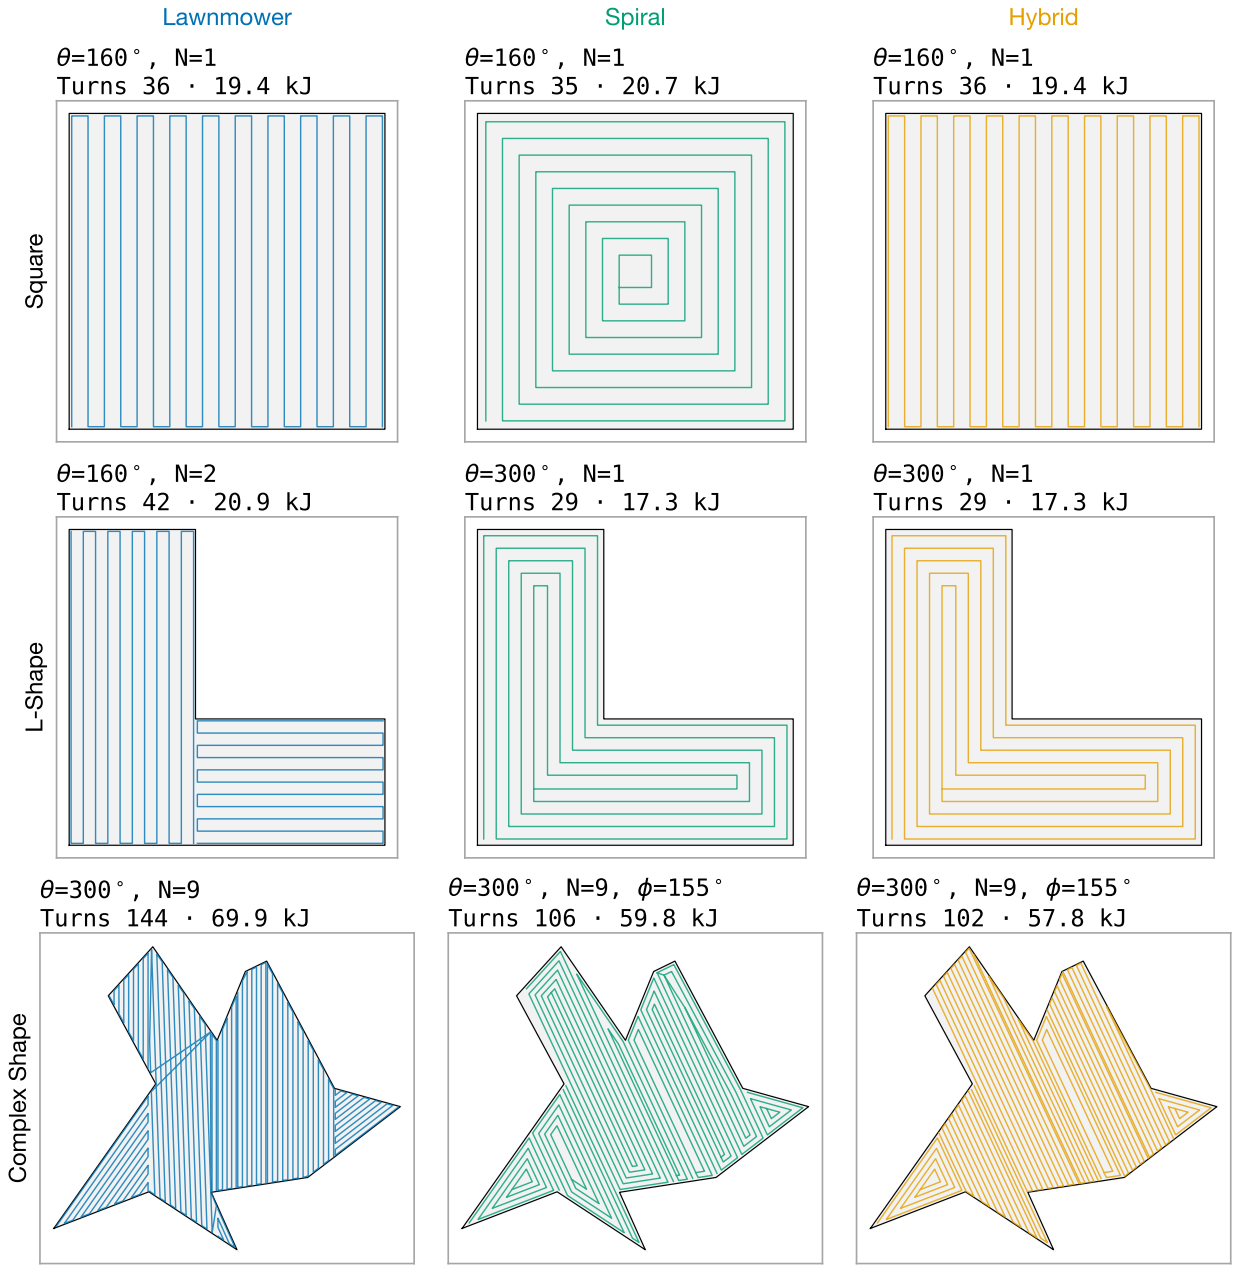


Figure 3. Visual comparison of lawnmower, spiral, and hybrid strategies on three polygons at the headline single-pass operating point ($\rho = 0$), drawn at a reduced sensor footprint (two-thirds of the benchmark value) so the sweep lines are dense enough to read; the hybrid uses the same per-cell selection as the baseline comparison (Section 4.1); the lawnmower column is drawn on the axis-aligned decomposition.

Table 2. Benchmark suite summary.

Category	Polygons	Area (m^2)	Compactness	Vertices	w (m)
Convex compact	9	43–260	0.60–1.00	3–100	0.50–1.00
Convex elongated	12	40–200	0.22–0.82	4–6	0.35–1.00
Mildly concave	11	40–226	0.21–0.51	6–24	0.50–0.80
Deeply concave	10	58–224	0.12–0.54	8–34	0.40–0.80
Multi-concavity	8	66–200	0.26–0.66	12–13	0.50–0.80
Total	50				

2.5.2. Grid Search Configuration

A comprehensive grid search was conducted over 50 benchmark polygon instances, sweeping 8 angle thresholds (160–360°) × 6 overlap factors (0.0–0.5) × 3 strategies, giving 7200 solver runs in total. We stress that these are not 7200 independent samples: the three strategies are the methods under test, so the number of independent problem instances is the 50 polygons, and all inference (Section 3) is performed at that level. A strategy is inapplicable at a configuration when one of its cells has no valid path. The lawnmower and the hybrid produced a plan at all 2400 configurations and the spiral at 2398: at the other 2, one cell's largest inscribed disc is narrower than the footprint, so no spiral ring fits it.

Unless stated otherwise, the per-strategy result figures (energy and coverage by category, the coverage distributions, computation time, and energy against coverage) display a single reference setting: no overlap ($\rho = 0$) and a moderate angle threshold of $\theta = 220°$ (chosen as a fixed reference for the per-strategy figures; the full threshold sweep is given in Figure 4, and the per-polygon search selects θ individually), with each cell swept along its minimum-width direction. We show $\rho = 0$ because almost every operating point lies there (all but three of the hybrid ones); increasing overlap adds cost to all three strategies (the full overlap sweep is given separately in Figure 5). Table 3 pools the grid configurations in which all three strategies yield a valid path (paired configurations), the per-category coverage means of Section 3.6 average each strategy over its own valid configurations, and the per-polygon comparisons (Wilcoxon test, dominance counts) use each strategy's operating point, treating the 50 polygons as the independent observations. The path-comparison figure (Figure 3) is rendered at the headline single-pass operating point ($\rho = 0$), matching the baseline comparison of Section 4.1; the coverage-depth heatmap (Section 3.6) uses $\rho = 0.5$ so that the double-coverage bands are dense enough to read. This is noted in their captions.

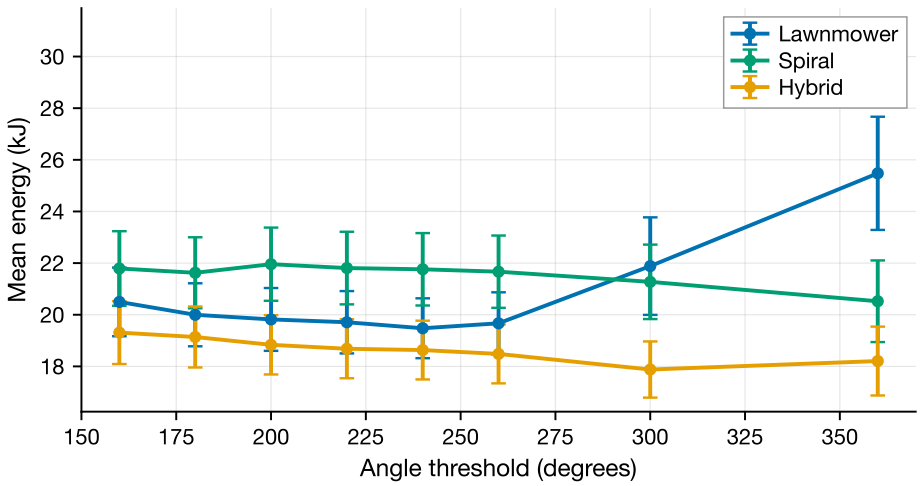


Figure 4. Effect of decomposition angle threshold on mean energy at $\rho = 0$.

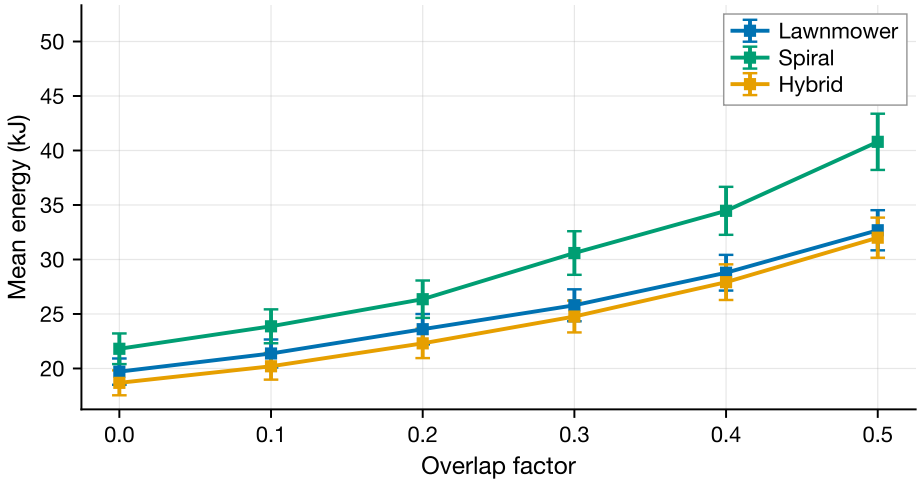


Figure 5. Effect of overlap factor on energy consumption.

Table 3. Strategy comparison pooled over the full configuration sweep. Mean ± SD.

Metric	Lawnmower	Spiral	Hybrid
Energy (kJ)	27.0 ± 14.9	29.4 ± 15.3	24.4 ± 11.7
Turns	51.9 ± 30.0	56.5 ± 34.2	45.3 ± 24.3
Coverage (%)	98.5 ± 1.5	98.6 ± 2.4	98.3 ± 2.3
Path length (m)	270 ± 149	258 ± 126	250 ± 119
Mission time (s)	176 ± 97	192 ± 101	159 ± 77

For each configuration, we record total energy E_{total} (J), turn count T, coverage (%), path length (m), mission time (s), and computation time (s). Coverage is the area of P covered by the union of the $w/2$-buffers of the path's straight segments, divided by the area of P; the union is formed segment by segment, because buffering a self-crossing path as one polyline can count unswept enclosed areas as covered. This assumes gimbal-stabilised nadir sensing at constant altitude over locally flat ground, so that the footprint is a ground strip of constant width w, with detection treated as binary inside it; terrain relief, motion smear at the turns, and soft-edged spray swaths are not modelled, so the reported coverage is an optimistic geometric bound. The intended applications—photogrammetric mapping, infrastructure inspection, and agricultural spraying—are endurance-limited rather than sensor-limited, and the sensor enters the planner only through w and ρ. Because benchmark polygon areas range from 40 to 260 m^2 (6.5× variation), cross-category comparisons use energy density (J/m^2) to normalise for polygon scale. Paired comparisons compare the same polygon across strategies, where area cancels, and remain valid without normalisation. Three quantities are distinguished: coverage completeness C, an output (the fraction of P swept at least once, constrained by the floor); the overlap factor ρ, an input that sets the spacing $\Delta = w(1-\rho)$ ($\rho = 0.5$ is 50% side overlap); and double-coverage completeness, an output (the fraction of P covered by two or more footprint segments; in the double-coverage study of Section 3.10, it includes the overlap of consecutive segments at turns, so it is an upper estimate of coverage by two separate passes).

3. Results

3.1. Strategy Comparison

We first compare the three strategies qualitatively and then in aggregate. Figure 3 shows representative lawnmower, spiral, and hybrid paths on three polygons, illustrating how the patterns differ in sweep structure and turn count before we quantify them across the full suite. On the multi-concavity Complex Shape, one lawnmower strand near a reflex

vertex departs from the clean back-and-forth pattern of the rest of the panel: rather than terminating a row at the concavity, the minimum-width sweep routes it through the cell interior to reach the far sub-region.

Each panel is rendered at or near that strategy's own energy-optimal decomposition (threshold θ, cell count N, and inclined orientation ϕ annotated per panel), e.g., on the L-shape, the hybrid covers the region undecomposed with a contour-following spiral (matching the spiral column), whereas the lawnmower's cheapest plan decomposes it. Each panel reports the turns and energy of the core planner's chained path (the selection criterion is energy, not turns): on the Square, for instance, the spiral has one fewer turn than the lawnmower but higher energy, so the hybrid—which minimises energy—selects the lawnmower. On every panel shown, the hybrid's energy is at or below both baselines. Paths are illustrative renders; the reported turns and energy are the core planner's values.

Table 3 presents grand-mean strategy performance pooled over the full configuration sweep (all decomposition thresholds and overlap factors); the per-polygon operating-point analysis follows below.

The hybrid achieves the lowest mean energy (24.4 kJ), a 16.8% reduction in grand-mean energy relative to spiral and 9.6% relative to lawnmower. The headline comparison, however, is made at each strategy's own operating point per polygon—the configuration a planner would actually deploy and the level at which the comparisons below are computed—where the hybrid reduces spiral energy by a median 4.8% per instance (mean 8.6%; the 11% sometimes quoted is the ratio of the two means, which overstates the typical paired reduction). Mean coverage at these operating points is 96.5% (hybrid), 96.7% (spiral), and 97.4% (lawnmower), so the hybrid's operating points cover slightly less on average, 0.2 points less than the spiral's and 0.9 points less than the lawnmower's; the L-shape saving in particular is taken mostly in coverage (Section 4). Any residual perimeter margin at zero overlap shrinks at a practical overlap ($\rho = 0.5$, Section 3.6); energy and coverage are plotted together for every polygon and strategy in Section 4.1. We caution that the grid comprises a within-instance sweep over decomposition, overlap, and orientation settings, not a set of independent samples: the number of independent observations is the 50 polygons, and all inference is performed at that level.

3.2. Comparisons Between Strategies

Only one of the three pairwise comparisons is open a priori: the hybrid is the cheapest of three plans that include the uniform lawnmower and spiral, so its rank is fixed by the dominance counts below, and an omnibus rank test would add nothing. Inference is conditional on the designed suite (Appendix A), whose 50 polygons span five geometric categories rather than being drawn at random.

Lawnmower and spiral do not dominate one another a priori, so this pair is the only one that carries an inferential test: a Wilcoxon signed-rank test finds no significant difference between them ($p = 0.47$; matched-pairs rank-biserial $r_{\mathrm{rb}} = 0.12$, a small effect, Table 4); lawnmower is significantly cheaper only on the convex-elongated category ($p = 0.0068$, or 0.034 after Bonferroni correction over the five categories), spiral tends to be cheaper on the mildly concave ones ($p = 0.054$, not significant), and elsewhere neither dominates. This is an informative null: with the spiral baseline using its best valid variant per cell, no fixed pattern is reliably cheaper across the suite—which is cheaper depends on geometry, with spiral winning on 22 of 50 polygons—so committing to a single pattern leaves energy on the table, and per-geometry selection (the hybrid) is exactly the response. For the two comparisons involving the hybrid, the sign of each difference is fixed by construction wherever the fixed pattern's operating point meets the floor (Section 2.4), so they are reported as dominance counts with the magnitude and distribution of the improvement

(Table 5). Against lawnmower, the hybrid strictly improves on 29 of 50 polygons (median strict reduction 12.6%, range up to 42.3%, on the 100-vertex Circle) and ties on the remaining 21, with no losses at the 90% floor; the wins comprise the 22 polygons where spiral is cheaper and the hybrid recovers it, and seven where it mixes patterns across cells, while the ties are where lawnmower is already the cheaper whole-polygon pattern. Against spiral, the hybrid strictly improves on 30 of 50 polygons (median strict reduction 10.9%, up to 56.8% on the Stadium) and ties on the other 20 (where the spiral baseline's best variant already matches the hybrid's per-cell spiral choice).

Table 4. Lawnmower versus spiral at the per-polygon operating points: Wilcoxon signed-rank test, the only comparison without a winner fixed by construction. p is two-sided; category rows are also given after Bonferroni correction over the five categories. r_{rb} is the matched-pairs rank-biserial correlation (positive: lawnmower cheaper).

Polygons	n_{pairs}	p	p (Bonferroni)	r_{rb}	Cheaper
All	50	0.47	—	0.12	n.s.
Convex compact	9	0.43	1.00	0.33	n.s.
Convex elongated	12	0.0068	0.034	0.85	lawnmower
Mildly concave	11	0.054	0.27	−0.67	n.s.
Deeply concave	10	0.63	1.00	0.20	n.s.
Multi-concavity	8	0.25	1.00	0.50	n.s.

Table 5. Dominance counts for the hybrid at the per-polygon operating points (90% floor). "Strict" counts strict improvements and "Ties" equal energy; "Median Strict Reduction" and "Range" are taken over the strict improvements only.

Relation	Strict/Ties	Median Strict Reduction	Range
Hybrid ≤ Lawnmower	29/21	12.6%	0.2–42.3%
Hybrid ≤ Spiral	30/20	10.9%	0.1–56.8%

Table 5 summarises how often the hybrid strictly improves on each fixed strategy and by how much.

Table 6 repeats the comparison at stricter floors. At every floor, the hybrid stays at or below both fixed patterns on every polygon and reaches the floor on all 50 polygons, while its mean operating-point energy rises from 17.5 kJ at 90% to 22.2 kJ at 99%. At 99%, the lawnmower reaches the floor on only 41 polygons, so part of that row compares the hybrid with best-effort lawnmower plans; on the 41 polygons where both reach the floor, the hybrid is cheaper on 23 and equal on 18. The comparison between lawnmower and spiral moves with the floor: lawnmower is cheaper on 34 polygons at 95% (Wilcoxon $p = 0.027$) and spiral on 30 at 99% ($p = 0.020$), although on the 41 polygons where both reach 99%, the two are level ($p = 0.42$).

Table 6. The comparison at four coverage floors, with the grid re-run at each floor and each strategy at its own operating point. "W/T/L" counts the polygons on which the hybrid is cheaper than, equal to, or dearer than the lawnmower (L) or the spiral (S); "Med." is the median per-polygon saving and "RoM" the ratio-of-means saving (%); "Best Eff." counts the polygons on which a strategy has no configuration at the floor, where its operating point is its highest-coverage configuration; "Cov." is the mean coverage at the operating points (H: hybrid).

Floor	W/T/L vs. L	Med./RoM	W/T/L vs. S	Med./RoM	Best Eff. L/S/H	Cov. L/S/H (%)
90%	29/21/0	2.7/7.4	30/20/0	4.8/11.1	0/0/0	97.4/96.7/96.5
95%	24/26/0	0.0/6.0	35/15/0	7.9/12.3	0/0/0	97.9/98.4/97.9
98%	25/25/0	0.6/13.0	33/17/0	5.5/12.8	0/0/0	98.8/99.2/98.9
99%	32/18/0	11.1/24.2	26/24/0	1.1/8.4	9/0/0	99.2/99.5/99.4

3.3. Effect of Angle Threshold and Overlap Factor

Having established the aggregate ranking, we now examine how the two solver parameters—the decomposition angle threshold θ and the overlap factor ρ—shape mean energy, since it is θ that the per-polygon search tunes. Figure 4 plots mean energy against θ.

At $\rho = 0$, mean lawnmower energy falls from $\theta = 160^\circ$, where polygons are split into many cells, to a minimum at 240° (19.5 kJ), and rises steeply at 300° and 360° (25.5 kJ), where concave polygons are swept as one or two large cells whose split rows must be joined by long in-polygon connectors. The spiral mean is nearly flat and lowest at 360°; the hybrid mean falls to its minimum at 300° and rises slightly at 360°. These marginal means do not show which threshold each polygon selects: the per-polygon search usually returns little to no decomposition and a multi-cell solution mainly on concave shapes. The fixed $\theta = 220^\circ$ used elsewhere is a representative moderate reference, not a global optimum.

Raising the overlap factor packs the sweep lines more densely and adds energy to every pattern—from $\rho = 0$ to 0.5 at $\theta = 220^\circ$, 66–71% for lawnmower and hybrid and 87% for spiral—without reordering the strategies (Figure 5).

3.4. Ablation of the Method's Components

Table 7 switches off one component of the method at a time and compares the resulting operating points with those of the full method at the 90% floor, all costed by the same planner and energy model.

Every component except the cell order saves energy on some polygons and costs it on none. The threshold search contributes most: fixing θ at 220° costs more on 20 polygons (ratio of means 6.1%), and at 180° on 25 (8.4%). The orientation search saves on 16 polygons, the multi-start tour on 10, and per-cell mixing over a per-polygon choice on 9; the fallback to a uniform plan changes one polygon. The exact order and direction of the cells, found by enumeration on up to six cells and by a Held–Karp dynamic programme on seven, would be cheaper on 8 of the 13 multi-cell operating points, where the multi-start tour costs at most 3.7% more. The variants that restrict the search cover slightly more on average (up to 97.3% against 96.5%), because restricting the search removes cheap configurations just above the floor. The components are removed one at a time, not in a factorial design, and the slab orientation is tuned on the mixed plan and shared, which can favour the hybrid over the per-polygon variant. At the stricter floors, per-cell mixing still saves energy over the per-polygon choice on 8–9 polygons and costs it on none.

Table 7. Ablation at the 90% floor, each variant against the full method's operating points: θ restricted to one value; slabs kept axis-aligned; the cheaper of the uniform-lawnmower and uniform-spiral operating points; the per-cell-mixed plan without the transit-inclusive fallback to a uniform plan; and the cells of each operating point re-chained by one nearest-neighbour tour or in the exact order. "Full/Tie/Variant" counts the polygons on which the full method is cheaper, the two are equal, or the variant is cheaper; "Median" and "RoM" are the full method's median per-polygon and ratio-of-means saving relative to the variant, $100\,(E_{\text{variant}} - E_{\text{full}})/E_{\text{variant}}$ (%); coverage is the variant's mean at its operating points (full method 96.5%).

Variant	Full/Tie/Variant	Median	RoM	Coverage (%)
θ fixed at 180°	25/25/0	0.6	8.4	97.3
θ fixed at 220°	20/30/0	0.0	6.1	97.0
Orientation search off	16/34/0	0.0	2.7	96.8
Per-polygon pattern choice	9/41/0	0.0	1.5	96.8
Mixing without fallback	1/49/0	0.0	0.1	96.5
One nearest-neighbour tour	10/40/0	0.0	1.3	96.6
Exact cell order	0/42/8	0.0	−0.4	96.5

3.5. Category-Dependent Performance

The aggregate ranking conceals strong variation across geometry; Figure 6 breaks energy and coverage down by polygon category and strategy.

In the category means over the whole grid, spiral is the most energy-expensive strategy in every category (2–24% above lawnmower), and the hybrid the cheapest; Figure 6 shows the same ordering at the reference setting. At the per-polygon operating point, however—the configuration a planner actually deploys—spiral is the lower-energy whole-polygon choice on 22 of the 50 instances (lawnmower on the other 28). These 22 spiral-favourable polygons are a subset of the 29 on which the hybrid strictly improves on lawnmower; the other 7 strict improvements arise from per-cell mixing, where within a single decomposition, the hybrid selects spiral on some cells and lawnmower on others. The spiral-favourable set includes compact convex shapes (e.g., the Circle N-gon and the equilateral triangle—single-cell polygons where a contour-following spiral incurs fewer expensive direction reversals than a boustrophedon sweep, in line with the energy advantage of spirals reported by Cabreira et al. [3]); 9 of the 11 mildly concave shapes, half of the deeply concave ones, and 2 of the 8 multi-concavity shapes, where a contour-following pass reduces wasted sweeps; and 3 elongated shapes. That the spiral remains cheaper on this many shapes even against a turn-minimising minimum-width lawnmower is what makes per-cell strategy selection worthwhile: at the 90% floor, lawnmower is cheaper on more polygons (28 of 50), though not significantly, and the hybrid recovers whichever pattern is cheaper, and its per-cell mix, without committing to one.

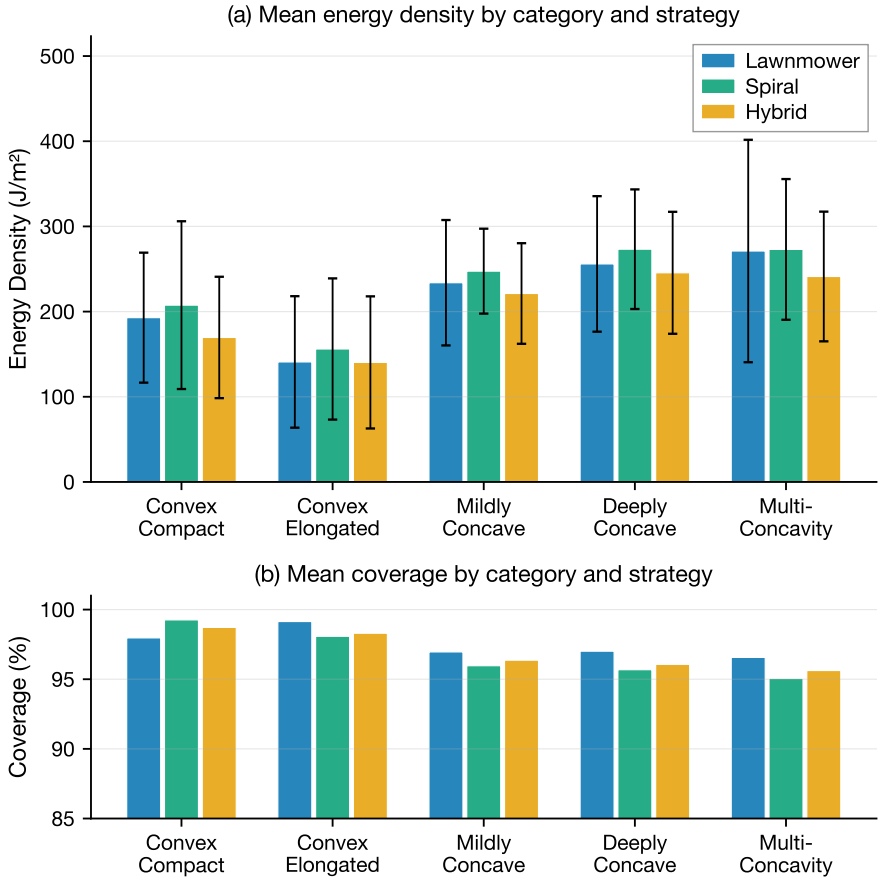


Figure 6. Energy and coverage by polygon category and strategy at the reference setting ($\theta = 220°$, $\rho = 0$): (**a**) mean energy density (J/m^2); (**b**) mean coverage (%). Spiral has the highest energy density in every category; its coverage is highest only on the convex-compact shapes, and lawnmower has the highest coverage on the other four categories.

3.6. Coverage Analysis

Energy is only one of the two co-objectives; a low-energy path is useless if it leaves the region under-observed. We therefore characterise how completely each strategy covers the domain, how that completeness depends on overlap, and where coverage defects concentrate. The distribution of coverage by strategy and by category is summarised in Figure 7.

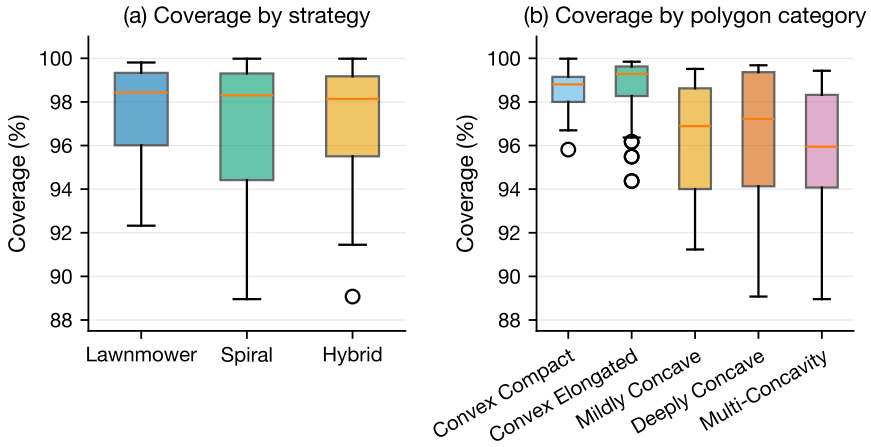


Figure 7. Coverage distribution by (**a**) strategy and (**b**) polygon category.

Where the boxplots summarise coverage as distributions, Figure 8 shows its spatial pattern, mapping coverage depth across two representative polygons.

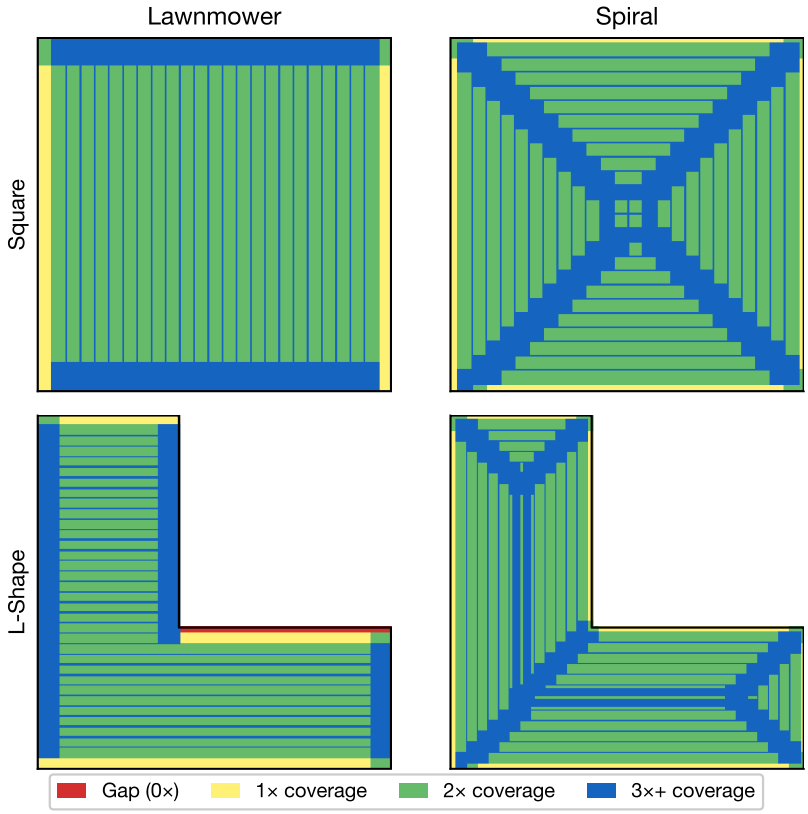


Figure 8. Coverage depth heatmaps for lawnmower and spiral strategies at overlap $\rho = 0.5$ (denser than the $\rho = 0$ operating point of Figure 3). Red = gap, yellow = single coverage, green = double, blue = triple+. At $\rho = 0.5$, both strategies achieve near-complete coverage, so the maps are dominated by double coverage; the residual single-coverage strips, and on the L-shape a narrow uncovered strip along the inner edge, lie at the polygon boundaries.

Coverage depends strongly on the overlap factor, and the strategies' relative coverage must be read at a fixed ρ. Averaged over the paired configurations, all three strategies cover almost the whole domain: spiral has the highest mean (98.6%), with lawnmower (98.5%), and hybrid (98.3%) marginally behind. Coverage rises with overlap, and the strategies converge, with spiral gaining a slight edge at high overlap because the contour-following rings close the perimeter margin (Figure 8). Coverage is high and roughly uniform across geometric categories (mean lawnmower coverage 97.3–99.4% across categories; Table 8), so the small residual is mainly a generic boundary-margin effect of boustrophedon sweeping—thin uncovered strips along the polygon edges where individual sweep lines terminate. The contour-following spiral closes even these margins at higher overlap, which is why it edges ahead by $\rho = 0.5$, while at low overlap, the lawnmower is already competitive. The residual coverage difference is a genuine co-objective rather than a confound to the energy comparison: where the spiral covers more than both other strategies (15 polygons at the reference setting), it does so at higher energy and, on all but one, with more turns than the lawnmower. Energy and coverage are plotted together in Section 4.1. The hybrid's energy is bounded above by the lawnmower's; it recovers the spiral where that is locally cheaper and tracks the lawnmower's coverage closely.

The largest residuals occur on complex and thin geometry. Figure 9 maps them for the lawnmower at the reference setting (97.5% covered on the Complex Shape benchmark of Li et al. [12] with nine cells, 93.8% on the star with seven, at the axis-aligned $\theta = 220°$ decomposition). The dominant artifact is the thin perimeter strip where the half-footprint inset leaves the boundary under-swept. There is no footprint-wide uncovered band along the internal decomposition seams: every cell is swept with a half-footprint (centreline) inset, so along most of each seam, the footprints of adjacent cells meet; narrower slivers remain where a seam passes through a reflex vertex (visible on the star). On the thinnest cells, the perimeter strip is a larger fraction of the area, which is where the contour-following spiral retains a slight coverage edge, at higher energy; the residual shrinks as the overlap factor rises.

Aggregated over the whole suite, these residuals leave every category above 95% mean coverage, as Table 8 details.

We now test how the comparisons respond to changes in the energy model and in the problem scale.

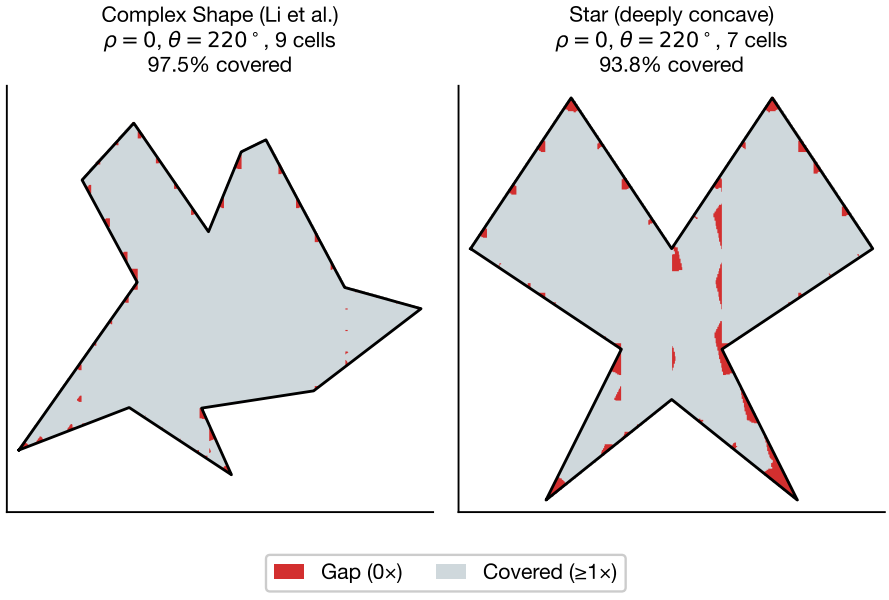


Figure 9. Coverage artifacts on complex geometry, rendered with the same per-cell oriented-boustrophedon generator as Figure 3. Lawnmower coverage of the multi-concavity Complex Shape

of Li et al. [12] (97.5% covered) and a deeply-concave star (93.8%) at overlap $\rho = 0$ and the reference threshold $\theta = 220°$, on the axis-aligned decomposition; red marks area not swept by the per-cell rows (inter-cell transits excluded), whereas the printed percentage is for the chained path. The residual is mainly the thin perimeter strip where the half-footprint inset leaves the boundary under-swept, much of it narrower than the drawn boundary line and hidden by it at this scale; narrower slivers also remain along internal seams that pass through reflex vertices (visible on the star).

Table 8. Mean coverage (%) by strategy and polygon category, averaged over all thresholds and overlap factors ($\rho = 0$–0.5) in the grid search, each strategy over its own valid configurations; these pooled means describe the sweep, not the deployed operating points.

Category	Lawnmower	Spiral	Hybrid
Convex compact	98.8	99.8	99.4
Convex elongated	99.4	99.0	98.9
Mildly concave	98.3	98.6	98.4
Deeply concave	98.1	98.0	97.7
Multi-concavity	97.3	97.5	96.7

3.7. Sensitivity Analysis

To test how the comparisons respond to the energy model, five scenarios (c_1, c_3, c_4, P_{hover}, mass) were each perturbed by ±20%; because cruise is 0.05% of the base-suite energy, the three cruise-curve coefficients act on almost nothing, the scenario labelled P_{hover} scales the in-place yaw energy ($P_{\text{hover}}\, t_{\text{yaw}}$), and the mass scenario scales the remaining non-cruise, non-yaw terms (lift over the acceleration and deceleration phases, kinetic work, and the avionics energy). For each of the 11 scenarios (baseline + 10 perturbations), energy was recomputed for each strategy's operating-point path using component-level scaling on the fixed paths. To additionally test robustness to the structure of the turn model (Section 2.3), a twelfth scenario replaced the primary multirotor stop–rotate–accelerate turn cost with the fixed-wing coordinated banked-turn ablation, evaluated exactly on the same paths, and a thirteenth removed the in-place yaw, as a vehicle with a gimbal-stabilised sensor could fly. Figure 10 reports the resulting sensitivity of the hybrid-versus-spiral saving for the parameter scenarios.

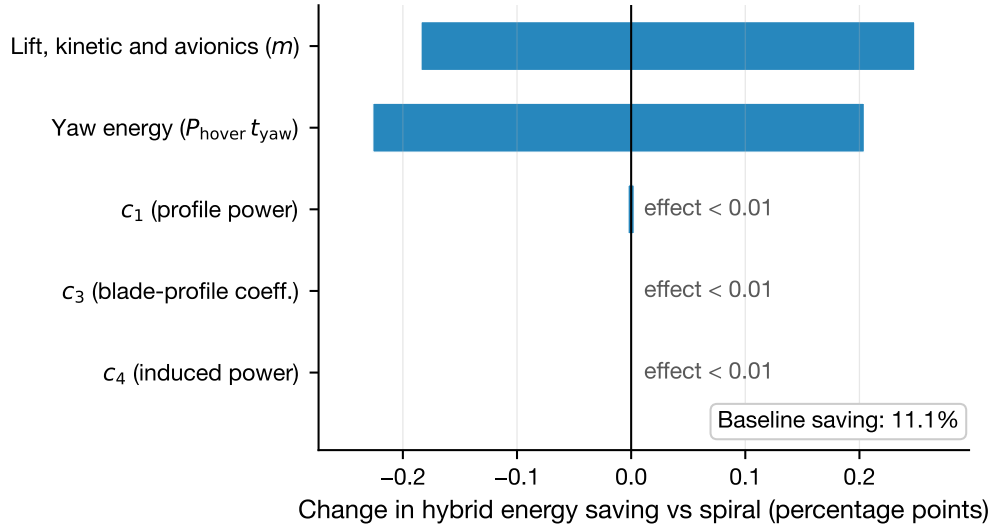


Figure 10. Hybrid-versus-spiral saving (ratio of means) under ±20% perturbation of c_1, c_3, c_4, the yaw energy, and the remaining non-cruise energy (lift, kinetic work, and avionics). The banked-turn ablation, not shown, lowers this saving to 3.9%, and removing the yaw lowers it to 9.5% (Section 3.7).

The hybrid had the lowest mean energy in all thirteen scenarios. Under the ten parameter perturbations, its ratio-of-means saving stayed at 10.9–11.4% against spiral and 7.0–7.9% against lawnmower (the corresponding paired per-instance reductions are lower, as in Section 3). Under the banked-turn ablation, the saving against spiral falls from 11.1% to 3.9%, while the saving against lawnmower rises slightly, from 7.4% to 7.7%; the two fixed patterns swap order in mean energy, and polygon by polygon, the hybrid is then dearer than the spiral on eight polygons and than the lawnmower on one. Without the in-place yaw, the savings fall to 9.5% against spiral and 4.5% against lawnmower, and the hybrid is dearer than the lawnmower on four polygons and than the spiral on one. The hybrid's lowest mean energy is therefore robust to the parameters and to the turn model, while its margins, and its freedom from losses, depend on the turn model.

3.8. Scale Sensitivity and Real-Boundary Demonstration

The base suite is turn-dominated, so we test how the results behave at larger scales, first as the region grows relative to the sensor footprint. When the first three polygons of each category are re-run at 1×–8× linear scale at fixed footprint width and at the single reference configuration ($\theta = 220°$, $\rho = 0$), with savings taken over the polygons on which all three strategies reach 90% coverage (Figure 11), the turn/maneuver share of lawnmower energy falls from ≈93% (mean area 94 m^2) to ≈74% (mean area ≈0.6 ha) as rows lengthen and true cruise phases appear. Across this range, the median hybrid-vs.-spiral saving rises from 7.0% to 13.3%, while the median hybrid-vs.-lawnmower saving falls from 1.9% to zero, because the lawnmower's own lead over the spiral grows with scale. Because the sweep scales row length and the region-to-footprint ratio together, it cannot attribute these changes to either alone.

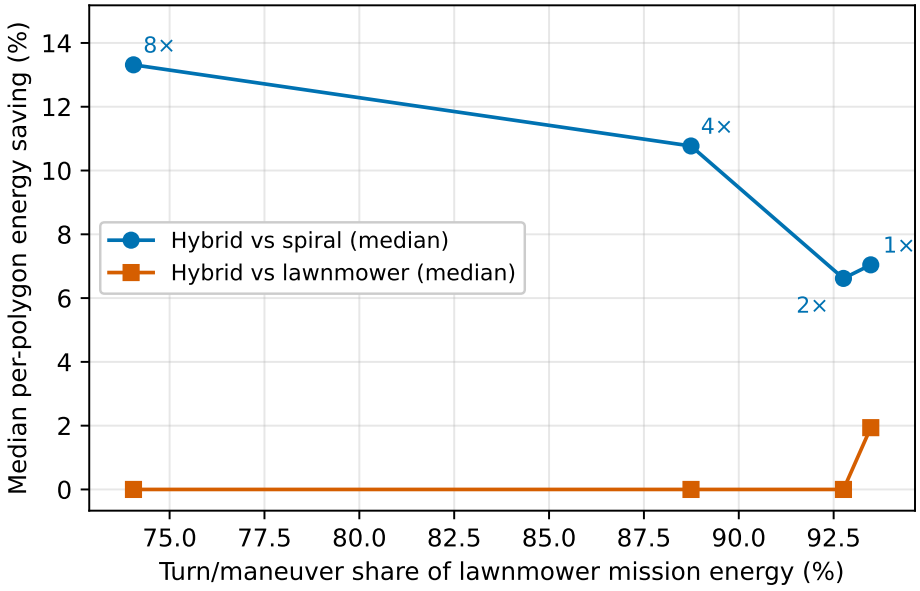


Figure 11. Median per-polygon hybrid saving versus the turn/maneuver share of lawnmower mission energy, over a 1×–8× linear scale sweep at fixed footprint width w and the single configuration $\theta = 220°$, $\rho = 0$ (annotations mark the scale factor). The hybrid-vs.-spiral saving persists across the regime; the hybrid-vs.-lawnmower gap closes as the hybrid's plan coincides with the lawnmower's on most polygons at larger scale.

Second, every polygon was scaled so that its footprint width, and hence its row spacing at $\rho = 0$, is 5 m or 10 m, which keeps A/w^2 and the polygon shapes but enlarges the fields to 0.3–1.8 ha and 1.2–7.1 ha, and the full grid, including the threshold, orientation, and

pattern searches, was re-run at cruise speeds of 5 and 10 m/s. The turn/maneuver share of lawnmower energy at the operating points then falls to 30–75%, with cruise taking 19–62%. At the 90% floor, the hybrid remains at or below both fixed patterns on every polygon in all four settings, with ratio-of-means savings of 6.1–7.2% against the lawnmower and 4.2–9.8% against the spiral (medians 2.1–6.6% and 0.0–5.2%), and mean coverage at the operating points of 96.1–96.7% for the hybrid (96.5–97.6% for the fixed patterns). At these spacings, the primary model overstates the turn cost, since it charges a full stop at every reversal; under the banked-turn model, whose 1.63 m turn radius (at 4 m/s and 1 g) fits between rows 5–10 m apart, the savings are 4.8–6.1% and 2.1–6.4%, and the hybrid is dearer than a fixed pattern on up to 11 polygons. Which fixed pattern is cheaper also depends on the setting: the spiral is cheaper than the lawnmower on 27 and 23 of the 50 polygons at 5 m spacing (5 and 10 m/s) and on 32 and 26 at 10 m spacing, significantly only at 10 m and 5 m/s (Wilcoxon $p = 0.020$).

To show the method solving the target real-world task end to end, we applied it to a precision-agriculture survey-mission problem: plan an energy-minimal coverage flight over a real farm field whose boundary is taken directly from a public GIS source. Six OpenStreetMap polygons (0.46–0.77 ha, 9–20 vertices, projected to metric UTM) were run through the identical core planner with the footprint width scaled to each field so that $A/w^2 = 80$ ($w = 7.56$–9.81 m), at the single configuration $\theta = 220°$, $\rho = 0$, with the orientation search but without the threshold and overlap search or the coverage floor (Figure 12). On one field, the chosen slab orientation leaves a cell whose largest inscribed disc is narrower than the footprint, so the spiral has no plan there. The hybrid reduces energy over spiral by a median 12.8% over the five fields where the spiral has a plan and over lawnmower by 1.2% over all six, and is at or below both baselines on every field by construction. Coverage is not equalised across strategies: the hybrid reaches 74.4–96.6% (mean 90.3%), and on the field where it saves 19.7% over the lawnmower, it covers 11.4 points less (74.4% against 85.8%), so that saving is taken in coverage. This is a demonstration on real boundaries rather than a large real-world benchmark; a systematic evaluation on a standard real-field dataset, such as the 350-field Fields2Benchmark [34], is left to future work.

We last report the computational cost.

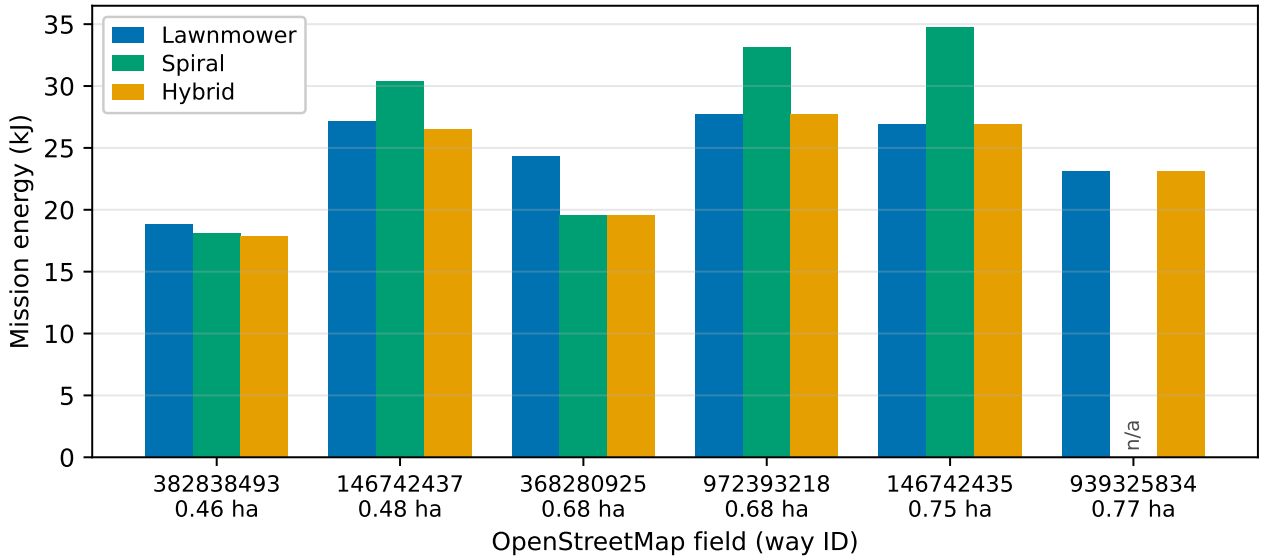


Figure 12. Planner on six real OpenStreetMap farmland field boundaries (single-pass $\rho = 0$, single configuration $\theta = 220°$, $A/w^2 = 80$): mission energy by strategy. The hybrid is at or below both baselines on every field; on the last field, the spiral has no valid plan at the chosen orientation (n/a).

3.9. Scalability

The per-configuration cost is dominated by the 5° decomposition-orientation search (36 slab orientations, each decomposed and cost-evaluated), which is shared across the three strategies at each grid point; a strategy-configuration then adds a per-cell minimum-width sweep and in-polygon connector routing. The mean cost is ≈530 ms per strategy-configuration (median 270 ms, maximum ≈ 5.6 s on the largest multi-cell shapes). Figure 13a plots computation time against sub-polygon count: the correlation is moderate ($r \approx 0.6$) because the orientation-search cost scales with cell complexity and the geodesic connector routing is invoked only on concave cells whose sweep lines split. The whole 7200-configuration grid completes in about eight minutes on eight cores. A complete plan for a new polygon, which evaluates all eight thresholds with their orientation searches at a given overlap factor, takes 13 s on average on one core (9 s median, up to 96 s on the 34-vertex Horseshoe); regenerating an already chosen configuration is sub-second.

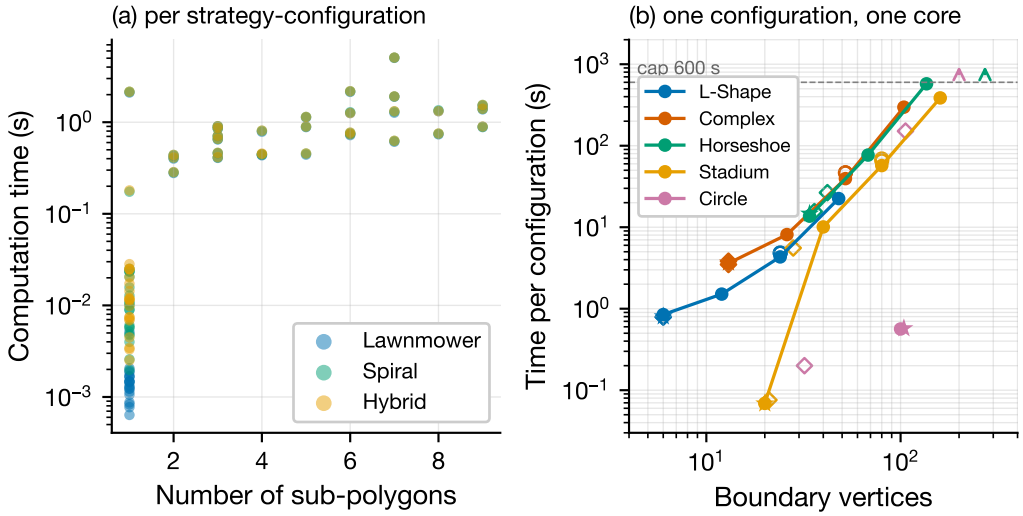


Figure 13. (**a**) Computation time vs. sub-polygon count (log scale), per strategy-configuration at the reference setting. Per-configuration cost is dominated by the 5° decomposition-orientation search; the moderate correlation ($r \approx 0.6$) reflects the cell count plus the geodesic connector routing used on concave cells. (**b**) One configuration timed on one core against the number of boundary vertices: densified boundaries (filled circles), jagged ones (open circles), the eightfold-densified boundaries after removing collinear points at $10^{-6}\,w$ (stars), the jagged ones after simplification at $0.1\,w$ (diamonds), and densified runs that exceeded the 600 s cap (arrows; the jagged and tolerance-zero runs that also exceeded it are not drawn).

The cost grows steeply with the number of boundary vertices. One configuration ($\theta = 220°$, $\rho = 0$, with the orientation search) was timed on one core for five boundaries densified two-, four-, and eightfold by inserting points along every edge, which leaves the shape unchanged, and with the inserted points displaced normal to their edge by up to $0.05\,w$ or $0.1\,w$ (Figure 13b). The merge dominates and grows roughly with the cube of the slab count: one configuration of the 34-vertex Horseshoe takes 14 s, at 68 vertices 77 s, and at 136 vertices 575 s, and at 272 vertices, it exceeds a 600 s cap. The plan also changes, because every vertex is a slab event: the fourfold L-shape is split into five cells instead of two at 28% more energy, and the eightfold Complex Shape into 20 instead of 7 at 21% more. Removing the inserted points first, by topology-preserving Douglas–Peucker simplification at a numerical tolerance of $10^{-6}\,w$, restored the original plan, to within rounding, on all five boundaries. Noisy boundaries need a larger tolerance, which on a finely digitised curve also removes original vertices and changes the plan (the densified Circle simplified at $0.05\,w$

keeps 38 of its 100 original vertices and costs 74% more). Boundaries should therefore be cleaned of collinear vertices before planning and simplified more coarsely only where they are noisy.

The single-pass results above all assume one coverage of the region; we turn to the complementary double-coverage setting, where a distinct pattern becomes advantageous.

3.10. Double-Coverage Patterns

Many UAV missions require double coverage (e.g., stereo photogrammetry with 60–80% side overlap). In the convention of Section 2.5, that percentage is ρ itself, so such missions are $\rho = 0.6$–0.8. The patterns below are instead compared at their sparsest double-covering settings—$\rho = 0.5$ for the two overlap-parameterised patterns, and $\rho = 0$ per pass for the cross-hatch and zigzag, which obtain their second look from a crossing pass—so this is the minimal double-coverage condition; denser overlap would raise all four costs. The established double-coverage pattern in structure-from-motion surveying is the cross-hatch (double-grid) flight, which improves reconstruction accuracy over a single grid at the cost of roughly doubling flight effort [35,36]. We therefore include cross-hatch as the principal incumbent baseline alongside boustrophedon 2×. Four strategies are evaluated (Figure 14):

1. Boustrophedon 2×: the minimum-width lawnmower of Section 2 with $\rho = 0.5$ (strips spaced at $w/2$).
2. Cross-hatch: Two perpendicular passes of that lawnmower at $\rho = 0$, the first along the minimum-width direction.
3. Boundary-reflective diagonal zigzag (novel): Continuous path bouncing between polygon boundaries at diagonal angles, advancing $w/2$ per half-bounce. All waypoints lie on the polygon boundary (zero external connection turns), and ascending/descending diagonals cross to provide implicit 2× coverage from two distinct viewing angles.
4. Spiral 2×: The continuous inward spiral with $\rho = 0.5$; the zigzag bounces along the longer side of the polygon's bounding box, and all four patterns are generated on the undecomposed convex polygon.

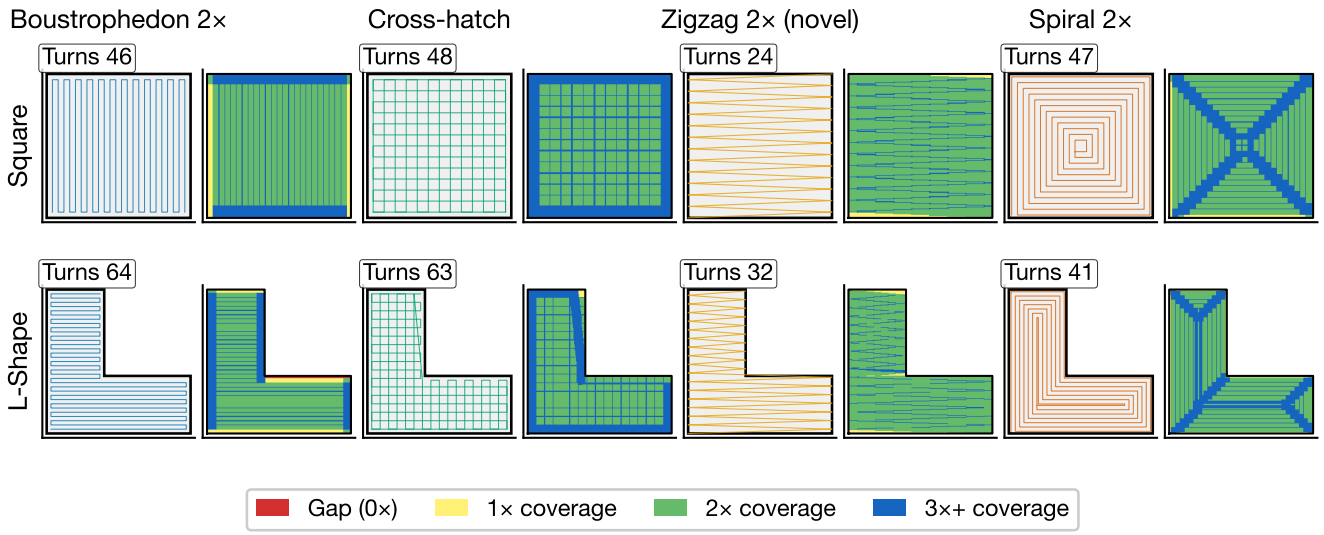


Figure 14. Double-coverage methods: flight paths and coverage heatmaps on Square and L-Shape; the lawnmower passes run along the minimum-width direction on the Square, as in Table 9, and axis-aligned on the L-Shape, which is shown for illustration only.

The turn, energy, and double-coverage figures behind these paths, averaged over the 21 convex polygons, are collected in Table 9.

Table 9. Double-coverage study (means across 21 convex polygons).

Method	Mean Turns	Energy (kJ)	2× Coverage (%)
Diagonal zigzag (novel)	23.5	20.1	93.8
Boustrophedon 2×	43.9	24.7	90.6
Spiral 2×	57.2	31.5	91.8
Cross-hatch	84.8	40.0	98.3

The boundary-reflective diagonal zigzag uses 46% fewer turns than boustrophedon 2× (23.5 vs. 43.9) and, under the stop–rotate turn model, 18% less energy (20.1 vs. 24.7 kJ), with higher double coverage on the metric of Table 9 (93.8% vs. 90.6%; that metric counts turn overlaps as a second pass, Section 2.5); on one of the 21 polygons, the Angled Corridor, boustrophedon 2× is cheaper. The advantage is larger against the cross-hatch incumbent: under the same model, the zigzag uses 72% fewer turns (23.5 vs. 84.8) and 50% less energy (20.1 vs. 40.0 kJ). These results are reported on the 21 convex polygons only: the zigzag is defined for convex and mildly concave cells whose two opposing boundaries are well defined, and strongly concave or multiply-connected regions must first be decomposed into such cells (Section 5). The comparison is also explicitly a turn/energy versus coverage-quality trade-off rather than a uniform improvement. Cross-hatch attains higher double coverage (98.3% vs. 93.8%) and—critically for structure-from-motion mapping—provides two orthogonal viewing directions, whereas the zigzag provides two oblique diagonal headings with a narrower angular baseline; viewing-angle diversity and near-complete double coverage are known to improve 3D reconstruction accuracy [35,37,38], so cross-hatch remains preferable when reconstruction quality, not survey effort, is the binding constraint. We also note that turn-cost-aware planners can minimise turns directly: Krupke [39] gives a near-optimal coverage planner with explicit turn costs. The zigzag is not optimal in that sense, and we do not attempt a head-to-head comparison here: Krupke's planner solves a combinatorial turn-cost optimisation, whereas the zigzag's appeal is precisely that it attains low turns and energy in closed form, with all waypoints on the boundary, without solving such a problem—so the two occupy different points on the optimality-versus-tractability spectrum and a fair comparison would require porting both to a common solver budget, which we leave to future work.

We do not claim that reflective (billiard-type) trajectories are themselves new—boundary-reflection dynamics are an established device in pursuit and patrolling problems [40], and stochastic "random-billiard" bounces have been used for probabilistic sweep coverage—but those are either patrolling (revisiting, not area completion) or randomised; applying a deterministic boundary-reflective sweep to achieve implicit double area coverage, with the half-bounce advance tied to the sensor footprint, is, to our knowledge, unreported in the UAV coverage literature. Taken together, the double-coverage results show the diagonal zigzag trading a modest coverage-diversity disadvantage for a large turn and energy saving (the latter under the stop–rotate turn model); we close the results with a rigid-body simulation check.

3.11. Independent Physics-Based Simulation Cross-Check

As a check that the planned paths and the model's speed profile can be tracked by a quadrotor, and that energy recovered from the rotor speeds preserves the ranking, we flew 150 paths (50 polygons × 3 strategies), the operating points of an earlier version of the method in which the corner where an inter-cell transit joins the next cell was not charged as a turn, with a 1.5 kg quadrotor under rigid-body dynamics in PyBullet [41], using the thrust/torque motor model of gym-pybullet-drones [42] and a geometric position/attitude controller, and recovered flight energy independently from the simulated per-rotor speeds; these are therefore not the operating points of Section 3. The commanded reference holds

the heading fixed and represents each in-place yaw by a hover dwell; the dwell and the analytical energy the flights are compared with use a yaw time proportional to the turn angle rather than the fixed t_{yaw} of Table 1. Two findings follow. First, with regard to tracking, the mean distance from the flown position to the nearest point of the planned path is 4.8 cm, a seventh to a twentieth of the footprint width (Figure 15). This measure does not register along-track lag or excursions over neighbouring rows: the largest deviations sit at the corners, where the vehicle lags the reference into each stop (on the 21 single-cell convex lawnmower paths, the 95th-percentile distance is 0.12–0.27 m), and the loops visible in Figure 15 occur where a transit joins the next cell, a corner these paths did not charge as a turn. Coverage is computed on the planned geometry and was not recomputed from the flown tracks. Second, with regard to energy consistency, the independently measured energy preserves the strategy ranking across the 150 flights, with the hybrid remaining at or below both baselines on every polygon, and differs from the analytical estimate only by a small, stable offset ($E_{sim}/E_{model} = 0.99 \pm 0.01$; Spearman $r_s = 0.999$, and at least 0.997 for each strategy). The commanded reference follows the model's own trapezoidal speed profile, so mission time agrees by construction. We do not claim absolute-Joule equivalence: the simulation reports rotor mechanical power (no battery or electrical-loss model), and its motor constants are matched to the analytical model, so the independent evidence is the trackability of that profile, the rotor-speed energy accumulated over it, and the preserved ranking, not the absolute power level. A hardware flight campaign remains the natural next step (Section 5).

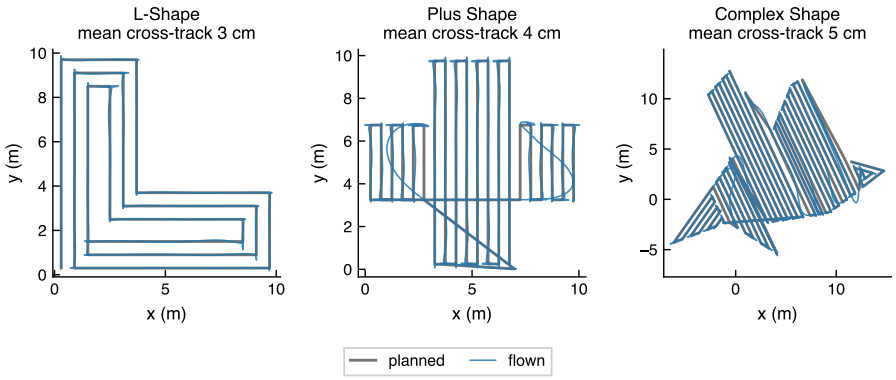


Figure 15. Rigid-body simulation in PyBullet: planned (grey) versus flown (blue) paths of a 1.5 kg quadrotor on three benchmark shapes. The annotation (labelled "cross-track") is the mean distance from the flown position to the nearest point of the planned path; this measure does not register along-track lag or excursions over neighbouring rows, which appear as loops in the Plus and Complex panels where a transit joins the next cell (a corner these paths did not charge as a turn).

4. Discussion

4.1. Comparison with Literature

Table 10 positions this work against representative CPP methods along the axes that distinguish them: whether the decomposition is energy-aware, whether more than one coverage pattern is available, whether decomposition granularity is itself searched under an energy or flight-time cost, how the per-cell pattern is selected, and the number of benchmark instances reported. The combination that characterises our contribution—energy-driven selection among multiple patterns and searched granularity, evaluated on a 50-polygon suite—does not appear in any single prior method: the two capability columns "Multi-Strategy" and "Granularity" are jointly populated by no prior row. The table covers methods that, like ours, decide coverage geometry offline; planners whose contribution is the online execution of a pattern fixed beforehand, such as Khaneghaei et al. [29], are not scored on these axes.

Beyond this feature-level positioning, Figure 16 puts the method head-to-head with three baselines on three shapes, and Table 11 with three planners on all 50 polygons, all costed by the same core planner and energy model.

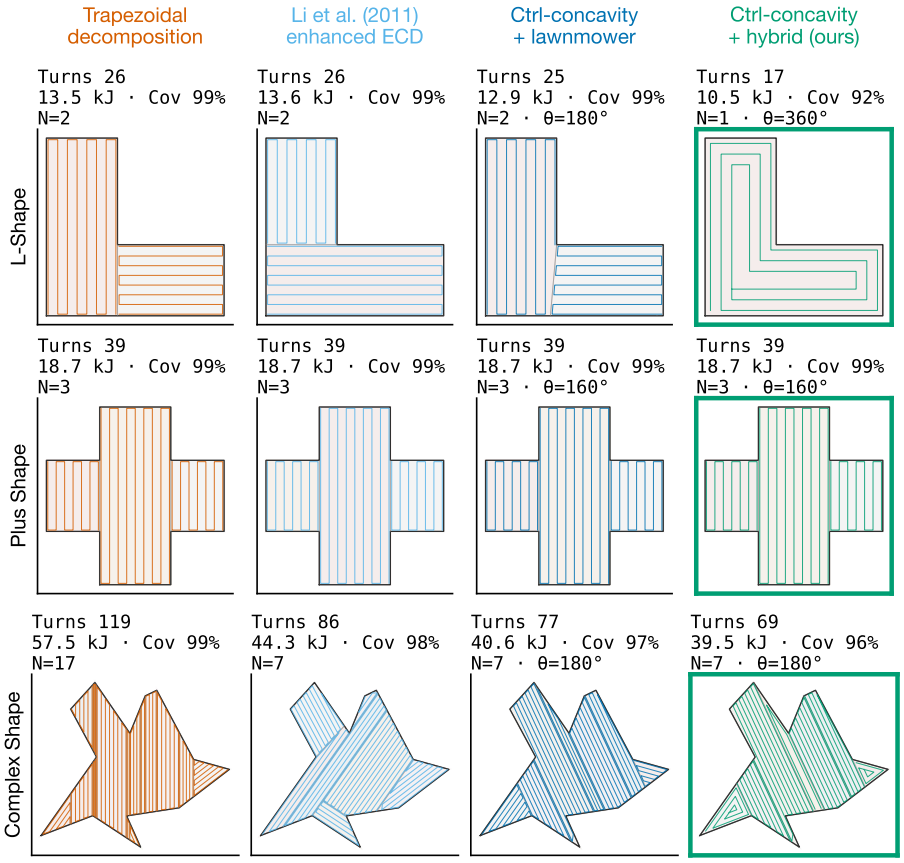


Figure 16. Four planners on three polygons at single-pass overlap $\rho = 0$. "Trapezoidal decomposition" slices the region at every vertex. "Li et al. (2011)" is our reimplementation of the minimum-width-sum decomposition of Li et al. [12] (greedy chord cuts through reflex vertices), each cell swept parallel to its minimum-width support line. "Ctrl-concavity + lawnmower" is the uniform lawnmower of this study at its own operating point for $\rho = 0$ (threshold searched, slab orientation shared with the hybrid). Every column is costed with the same per-cell minimum-width lawnmower and nearest-neighbour cell tour, except "ours", which shows the per-cell-mixed plan; coverage is printed on every panel. The green frame marks the lowest-energy planner among those reaching the 90% coverage floor; on the Plus, where all four tie, it is drawn on ours.

Table 10. Comparison with published CPP methods. "Multi-Strategy" = selects among more than one coverage pattern; "Granularity" = decomposition granularity is searched under an energy or flight-time cost (a turn-count or width-sum proxy, or a fixed, user-set tolerance, does not count).

Method	Decomposition	Energy	Multi-Strategy	Granularity	Selection	*N*
Choset [2]	Boustrophedon	–	–	–	–	5
Huang [11]	Line-sweep	–	–	–	–	3
Li et al. [12]	Enhanced ECD	–	–	–	–	1
Lien & Amato [6]	ACD	–	–	–	–	–
Brown & Waslander [5]	Constriction	–	–	–	–	–
Coombes et al. [10]	Convex (cost-searched)	Flight time (wind)	–	✓	–	33
Cabreira et al. [3]	–	Physics	–	–	–	3500+
Theile et al. [24]	Grid	Battery	–	–	Learned	3
Jia et al. [8]	Convex (concavity-elim.)	–	–	–	–	–
This work	BCD+merge (θ)	Physics	✓	✓	Energy	50

Table 11. The hybrid against three planners on all 50 polygons at four coverage floors (same polygons, footprint, energy model, coverage measure, and operating-point rule). Each cell gives the polygons on which the hybrid is cheaper/equal/dearer, and then its median per-polygon saving (%). "RoM" is the ratio-of-means saving against Fields2Cover; "best effort" counts the polygons on which a planner has no configuration at the floor; trap. = trapezoidal, F2C = Fields2Cover. At 99%, part of the margins compares best-effort plans.

	90%	**95%**	**98%**	**99%**
Li et al. (reimpl.)	34/16/0, 11.5	31/18/1, 9.2	30/18/2, 9.2	30/18/2, 11.8
Trapezoidal (reimpl.)	43/7/0, 22.7	42/8/0, 20.3	41/9/0, 20.5	40/8/2, 19.1
Fields2Cover v2.1	41/0/9, 6.9	41/0/9, 5.6	33/0/17, 2.9	29/0/21, 1.3
RoM vs. Fields2Cover (%)	9.5	6.7	1.1	−7.5
Coverage (%): hybrid/Li	96.5/98.2	97.9/98.4	98.9/98.8	99.4/99.0
Coverage (%): trap./F2C	98.6/98.3	98.6/98.8	98.9/99.4	99.0/99.7
Best effort: hybrid/Li	0/0	0/0	0/4	0/15
Best effort: trap./F2C	0/0	0/0	3/0	18/0

The proposed method selects the decomposition threshold and the per-cell strategy by the energy criterion, with the undecomposed single-sweep ($\theta = 360^\circ$) included as a candidate (a naive no-decomposition sweep is not shown separately: on convex regions it coincides with this candidate, and on the concave shapes here, an undecomposed lawnmower is valid but dearer than our plans: 26.8, 32.5, and 53.9 kJ on the L-Shape, Plus, and Complex Shape). Each panel reports energy, turns, and coverage: the comparison is not iso-coverage, because every method leaves a thin perimeter margin and the panels differ in coverage by up to seven points at $\rho = 0$ (Section 3). The green frame therefore marks the lowest-energy method among those reaching the 90% coverage floor, not the globally cheapest panel; on the Plus, where all four columns tie, it is drawn on ours. Paths are rendered with inside-the-cell helpers for legibility (every connector stays within the polygon), while the reported energy, turns, and coverage are the core planner's values on the same cells.

On the L-Shape, our search returns an undecomposed contour-following spiral (10.5 kJ), 19–23% cheaper than every lawnmower decomposition, Li et al.'s included (12.9–13.6 kJ), but at 92% coverage against their 99%, so the saving is taken mostly in coverage. At the 95% floor, the hybrid is 5% cheaper than Li et al., and at 99%, it is 0.2% cheaper. On the Complex Shape, the undecomposed lawnmower and spiral are valid but dearer (53.9 and 43.0 kJ), and our decomposition ($N = 7$) reaches 39.5 kJ at 96% against Li et al.'s 44.3 kJ at 98%, 11% cheaper at two points less coverage, mostly owing to the searched decomposition (a uniform lawnmower on the same cells costs 40.6 kJ). On the Plus, all four columns produce Li et al.'s convex cells ($N = 3$) and tie at 18.7 kJ and 99%.

Table 11 extends the comparison to all 50 polygons and the four floors. Li et al.'s enhanced ECD and the trapezoidal decomposition are our reimplementations, each cell swept by the minimum-width lawnmower and chained by the multi-start tour; the trapezoidal decomposition cuts vertical slabs at every vertex, with no merge and no orientation search, so it also splits 15 of the 21 convex polygons. Fields2Cover v2.1 [30] is independently authored and run with its boustrophedon decomposition (split angle searched in 15° steps), and brute-force swath angle and route planner, with no headland; our energy model charges every heading change of more than 10° along its route as a stop, as for our own paths. At the 90% floor, the hybrid is cheaper than Li et al. on 34 polygons (median 11.5%), than the trapezoidal decomposition on 43 (22.7%), and than Fields2Cover on 41, with Fields2Cover being cheaper on 9. The hybrid's operating points cover less, 96.5% against 98.2–98.6%. The margin over Fields2Cover narrows as the floor rises: Fields2Cover runs its swaths to the field boundary, without the half-footprint inset at the row ends of our patterns, and reaches 99% on every polygon; it is level with the hybrid on average at 98% (ratio of means

1.1%) and cheaper at 99%, where the hybrid's mean energy is 7.5% higher. At 99%, Li et al. and the trapezoidal decomposition reach the floor on only 35 and 32 polygons; where both planners reach it, the hybrid and Li et al. are level at the median, while the lead over the trapezoidal decomposition holds (median 14.1%). Figure 17 plots energy against coverage for every polygon and strategy at the reference setting.

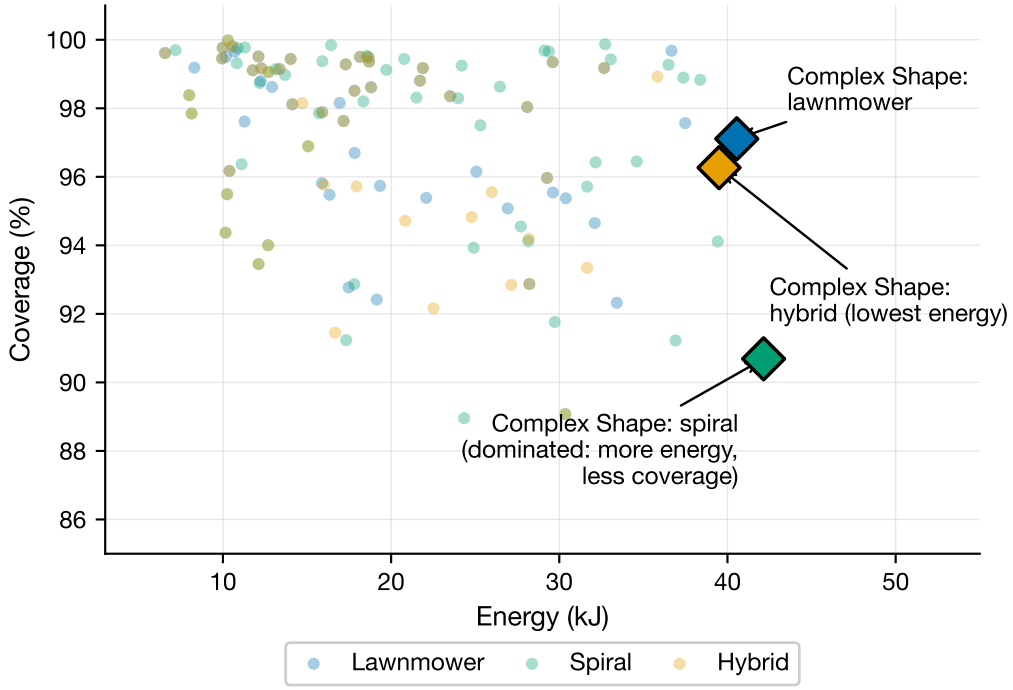


Figure 17. Energy against coverage for every polygon and strategy at the reference setting ($\theta = 220°$, $\rho = 0$); the three strategies on the Complex Shape are highlighted.

Relative to prior work, the contribution is not any single component—concavity-limited decomposition [5], energy-aware spiral planning [3], and physics-based multirotor power models [7] all predate this paper—but their integration into a single energy-minimising formulation in which decomposition granularity and per-cell pattern are selected by one energy criterion (with a minimum-width sweep), evaluated on 50 polygons with an analysis that is explicit about which comparisons are stochastic and which hold by construction.

4.2. Deployment Workflow

The method is designed to slot into a standard UAV survey pipeline as an offline planning stage. The end-to-end workflow is: (1) obtain the region boundary—either from a GIS/OpenStreetMap layer (as in the six-field demonstration of Section 3.8, polygons projected to metric UTM) or from an operator-drawn polygon in mission-planning software; (2) run the controlled-concavity decomposition and energy-driven per-cell strategy selection of Sections 2.1–2.4 for the vehicle's sensor footprint and required overlap, yielding the lowest-energy decomposition, per-cell patterns, and waypoint sequence among the evaluated candidates; (3) export the resulting waypoint list as a standard ground-control-station mission (for example a QGroundControl/MAVLink survey plan, or the equivalent waypoint mission for a Phantom-4-class platform) for upload and autonomous execution. Because a complete plan for a given boundary and overlap is computed in 13 s on one core on average for polygons of the size tested here (up to 96 s on the 34-vertex Horseshoe; Section 3.9), the planning stage can be re-run in the field if the boundary is revised; boundaries digitised with many short edges should first be cleaned of collinear vertices. The energy model is

calibrated to a ≈1.5 kg Phantom-4-class quadrotor (Section 2.3), so the parameters and the exported mission target a widely deployed survey vehicle rather than an abstract platform.

5. Conclusions

This paper presented an energy-aware 2D CPP framework combining a controlled-concavity boustrophedon decomposition with per-cell strategy selection. The key findings are:

1. At the per-polygon operating point (90% coverage floor), the hybrid strategy reduces spiral energy by a median 4.8% per instance (mean 8.6%; 11% as a ratio of means, 3.9% under the banked-turn model), at operating points that cover 0.2 points less than the spiral's on average. Lawnmower and spiral are not separable in aggregate (Wilcoxon $p = 0.47$; lawnmower wins significantly only on the convex-elongated polygons); the hybrid is at or below the lawnmower baseline by construction at every configuration—strictly improving on 29/50 shapes at the operating points, tying on the rest, and never losing at that floor—rather than by a statistical margin. The hybrid is the cheapest strategy in every geometric category; on all 50 polygons at the same floor, it is cheaper than our reimplementations of Li et al.'s enhanced ECD and of a trapezoidal decomposition on 34 and 43 polygons, with no losses, and than Fields2Cover on 41 (dearer on 9), its operating points covering 1.7–2.1 points less on average.
2. At 95–99% floors, the hybrid remains at or below both fixed patterns on every polygon and reaches each floor on all 50 polygons, while its mean energy rises by 27% from the 90% to the 99% floor; Fields2Cover, whose rows reach the field boundary, is level with it on average at 98% and cheaper at 99%.
3. Treating decomposition granularity as a search variable, rather than fixing it, lets the search favour little to no decomposition on most shapes and decompose mainly the concave ones; in the ablation, fixing the threshold at 220° costs more energy on 20 polygons and disabling the orientation search costs more on 16, while for the lawnmower, an intermediate threshold is cheapest on average.
4. The boundary-reflective diagonal zigzag achieves 46% fewer turns and 18% less energy than boustrophedon 2× for double coverage on the 21 convex polygons, at higher double-coverage completeness on the metric of Table 9 (93.8% vs. 90.6%; the higher-completeness cross-hatch retains a coverage and view-diversity advantage).
5. The hybrid keeps the lowest mean energy under ±20% energy-model parameter perturbation, under the banked-turn model, and without the in-place yaw, but its margins depend on the turn model (Section 3.7): under the two alternative models, it is dearer than the spiral on up to eight polygons and than the lawnmower on up to four. Over a 1×–8× scale sweep, the median saving against spiral grows while that against lawnmower falls to zero. With the polygons scaled to row spacings of 5–10 m (fields of up to 7.1 ha), where cruise takes 19–62% of the energy, the hybrid remains at or below both fixed patterns under the primary model, with ratio-of-means savings of 6–7% against the lawnmower and 4–10% against the spiral (medians 2–7% and 0–5%); under the banked-turn model, which suits these spacings better, it keeps the lowest mean energy but is dearer than a fixed pattern on up to 11 polygons.

Limitations

Several limitations bound the scope of these conclusions:

- The evaluation is entirely in simulation; we report no real-world (field) flight experiments. What we claim is accordingly a relative, model-based result: for every strategy, the reported energy is the cost of the planned path under one common power model based on published flight-validated multirotor power models, so the comparison iso-

lates the effect of the decomposition and pattern choices from the flight hardware. This ranking—rather than any absolute Joule figure on a specific vehicle—is the paper's claim, and in place of physical measurement we support it by (i) basing the power model on published flight-validated multirotor models (Section 2.3); (ii) showing that the hybrid keeps the lowest mean energy under ±20% parameter perturbation and two alternative turn models (Section 3.7), while a 1×–8× scale sweep shows how the margins change with scale (Section 3.8); and (iii) flying 150 paths of an earlier version of the method in a rigid-body physics simulation that recovers energy from the simulated motor speeds and preserves their strategy ranking (Section 3.11). A physical flight campaign measuring flown energy on a Phantom-4-class vehicle would additionally validate the absolute figures and calibrate the model to a specific platform; because it would test the power model rather than the planner's decisions, we treat it as complementary to—not a prerequisite for—the relative results reported here, and set it as the natural next step.

- The energy-model coefficients follow flight-validated multirotor power models [4,7], but the model itself is not directly flight-validated on the survey vehicle. Because all strategies are evaluated under identical parameters, relative comparisons are expected to be reliable in still air.
- The model assumes still air, a constant payload, and an ideal battery. In the base suite, payload and battery state act, to first order, as a common scaling, because about nine-tenths of each path's energy is drawn at the single hover power; at the 5–10 m spacings of Section 3.8, where cruise takes 19–62%, this argument does not apply, and their effect is not evaluated. Wind, the main source of direction-dependent leg power, does not act as a common scaling: it makes leg costs depend on heading, pulls the cheapest orientation and sweep towards the wind axis, and can reorder rankings (for fixed-wing surveys, wind-aware decomposition and sweep directions shortened flights by roughly 5–25% [10]). In the base suite, wind cannot be represented simply by evaluating Equation (4) on airspeed, because cruise is 0.05% of the pooled energy; at the spacings of Section 3.8, such a term would act on the 19–62% cruise share but is not evaluated, and a meaningful wind study also needs a power model in which the maneuver phases depend on airspeed. Gusts, turbulence and the extra thrust needed to tilt during acceleration are not modelled either.
- The method was evaluated from the turn-dominated base suite, in which the turn/maneuver share is ≈93% of mission energy, to fields of up to 7.1 ha at row spacings of 5–10 m, in which it is 30–75% (Section 3.8). Within this range, the hybrid stays at or below both fixed patterns under the primary turn model, but the margins depend on the setting: in the fixed-footprint scale sweep, at a single reference configuration, the median saving against the lawnmower falls to zero, while at 5–10 m spacing with the full grid, it is 2.1–6.6% (6.1–7.2% as a ratio of means). At such spacings, the primary model overstates turn cost, and under the banked-turn model, the hybrid is dearer than a fixed pattern on up to 11 polygons; larger fields with longer rows are untested.
- Benchmark polygons are predominantly synthetic and form a designed suite rather than a random sample, so the comparisons are conditional on it; we add a small demonstration on six real OpenStreetMap farmland boundaries (Section 3.8), but a randomised or real-field evaluation in the spirit of Oksanen and Visala [14], for example on a standard real-field dataset (the 350-field Fields2Benchmark [34]), would further strengthen generalisability.
- The boundary-reflective zigzag is designed for convex polygons and mildly concave shapes whose two opposing boundaries remain well-defined (e.g., the L-shape in

Figure 14); strongly concave or multiply-connected regions require prior decomposition into such cells.

- The per-vertex 10° turn test makes energy depend on how curves are digitised (an arc drawn at about 20° per vertex pays a stop at every vertex, while the same arc drawn more finely stops at none); this is behind the largest savings in Table 5, on the Stadium against spiral and on the 100-vertex Circle against lawnmower, and a curvature-based turn model would remove it.
- Survey specifications often require 99–100% coverage, and approaching it is costly: the hybrid's mean operating-point energy rises by 27% from the 90% to the 99% floor, and at 99%, the lawnmower reaches the floor on only 41 polygons, its main residual being the strip its half-footprint inset leaves at the row ends (Table 6); Fields2Cover, whose swaths run to the field boundary, reaches 99% on every polygon (Section 4). Boundary passes or row-end extensions that would recover the last points of coverage are not analysed here.

Future work will extend the decomposition and selection framework to 3D mesh coverage (building facade inspection), extend the rows to the boundary, test on a randomised or real-field polygon set, and validate the planned-versus-flown energy and coverage on a Phantom-4-class quadrotor in a field trial, closing the loop between the power model and measured battery draw.

Author Contributions: Conceptualisation, A.T.; methodology, A.T.; software, A.T.; validation, I.I.; formal analysis, A.T. and V.M.; investigation, A.T.; data curation, A.T. and O.M.; writing—original draft preparation, A.T.; writing—review and editing, A.T., I.I., V.B., V.M., and O.M.; visualisation, A.T. and I.I.; supervision, V.B.; project administration, V.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by a University of Warwick PhD scholarship.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The benchmark polygons are drawn in Appendix A; their coordinates, the planner code, and the result tables are available from the corresponding author on request.

Acknowledgments: The authors thank the University of Warwick for providing computational resources for the grid-search simulations.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ACD	Approximate Convex Decomposition
BCD	Boustrophedon Cellular Decomposition
CPP	Coverage Path Planning
ECD	Exact Cellular Decomposition
MILP	Mixed-Integer Linear Programming
NN	Nearest-Neighbour
RTK	Real-Time Kinematic
SD	Standard Deviation
UAV	Unmanned Aerial Vehicle
UTM	Universal Transverse Mercator

Appendix A. Benchmark Suite

Figures A1 and A2 draw all 50 benchmark polygons with their area, vertex count, footprint width w, and region-to-footprint ratio A/w^2.

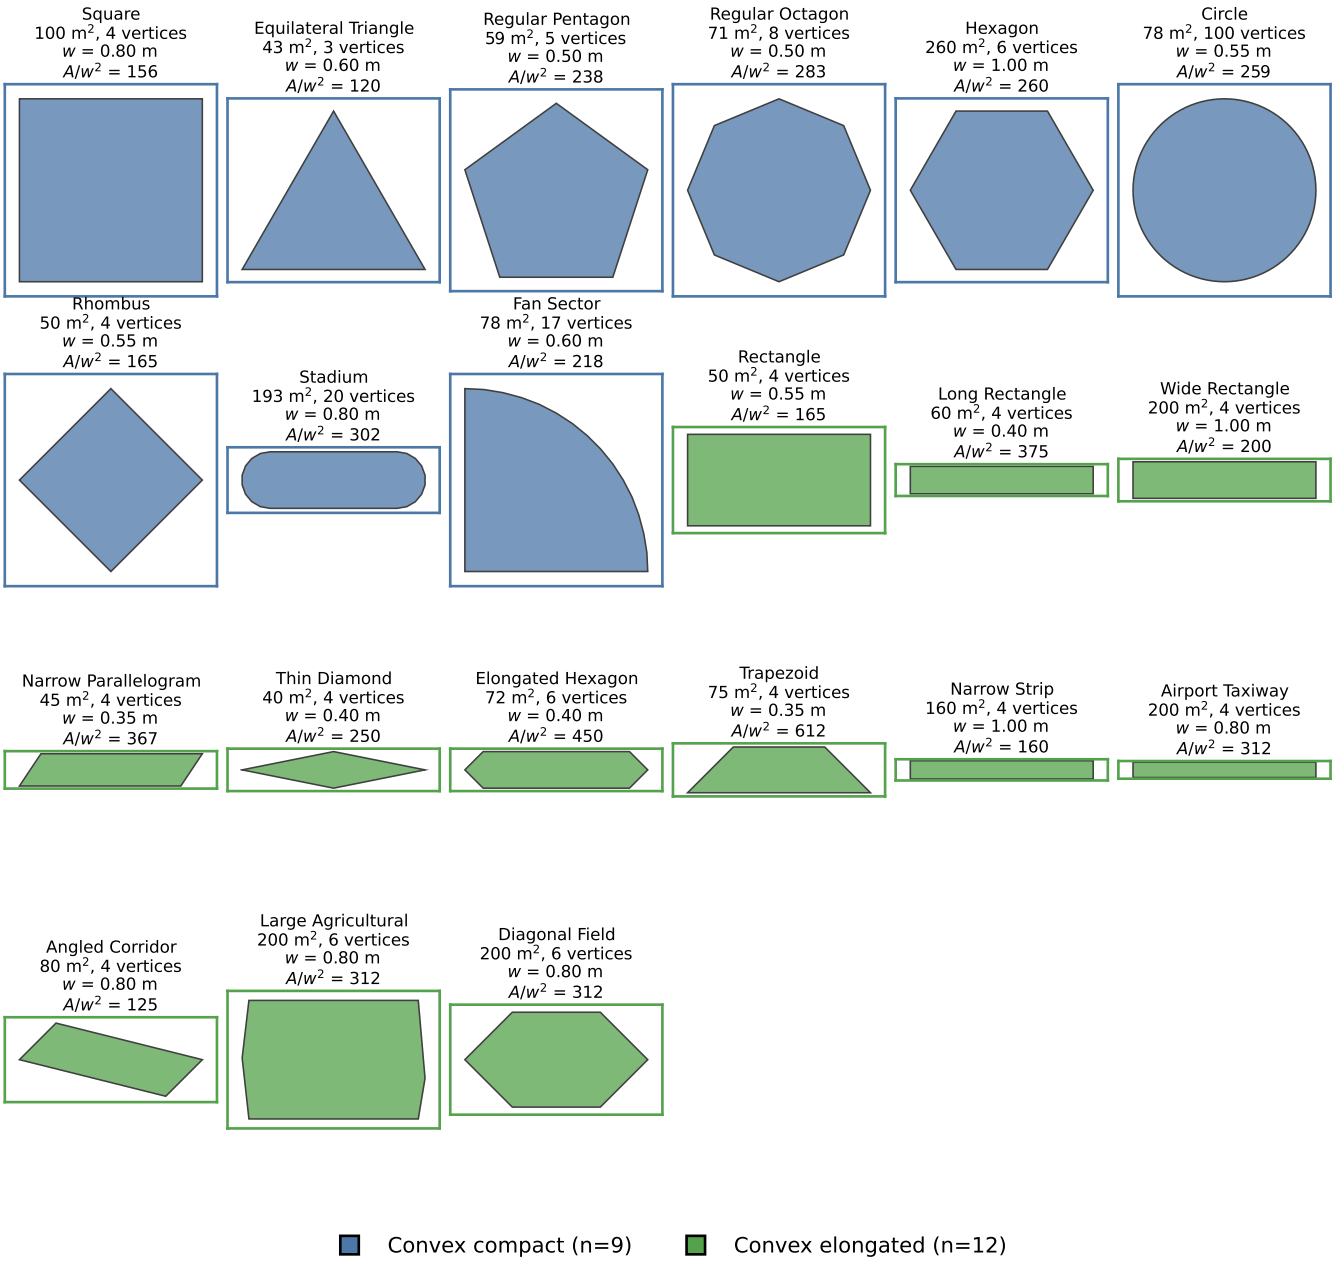


Figure A1. The 21 convex benchmark polygons, coloured by category, with area, vertex count, footprint width w, and A/w^2.

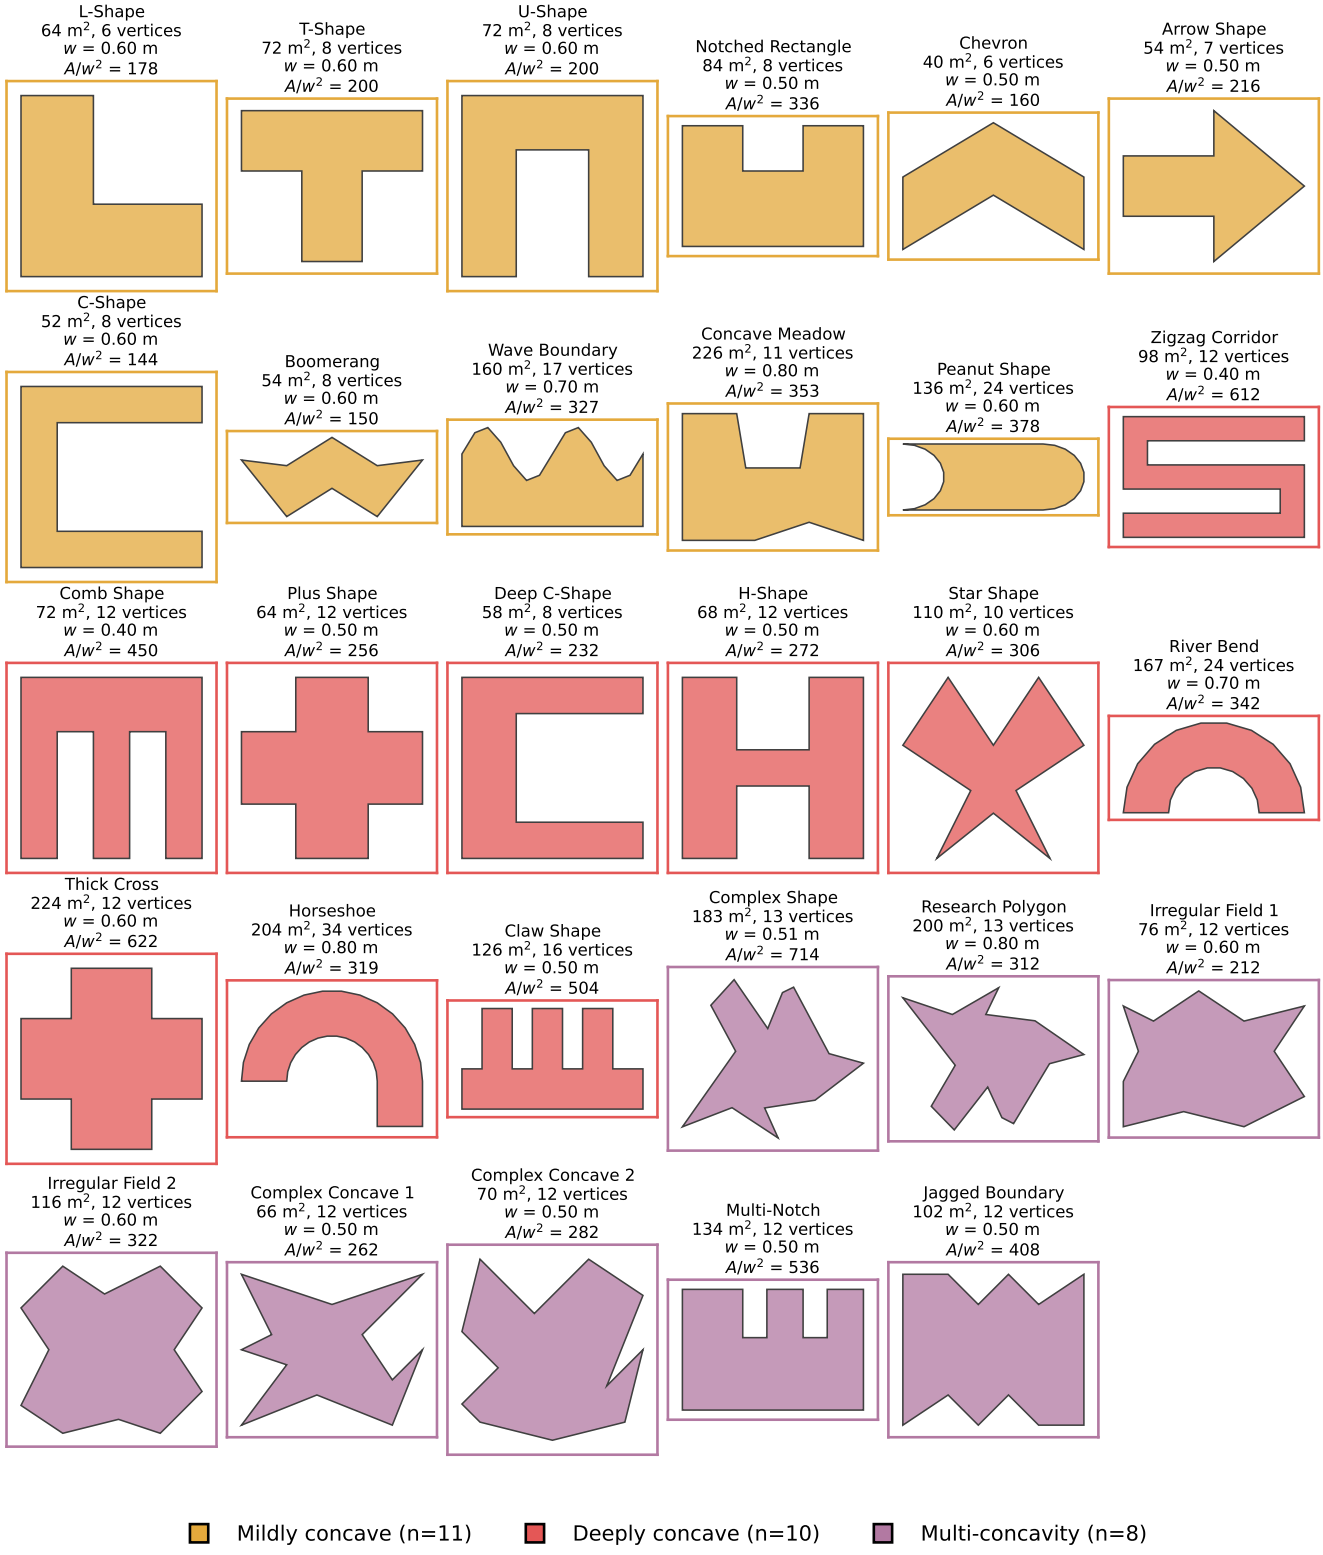


Figure A2. The 29 non-convex benchmark polygons, coloured by category, with area, vertex count, footprint width w, and A/w^2.

References

1. Galceran, E.; Carreras, M. A survey on coverage path planning for robotics. *Robot. Auton. Syst.* **2013**, *61*, 1258–1276. [CrossRef]
2. Choset, H. Coverage of known spaces: The boustrophedon cellular decomposition. *Auton. Robot.* **2000**, *9*, 247–253. [CrossRef]

3. Cabreira, T.M.; Di Franco, C.; Ferreira, P.R.; Buttazzo, G.C. Energy-aware spiral coverage path planning for UAV photogrammetric applications. *IEEE Robot. Autom. Lett.* **2018**, *3*, 3662–3668. [CrossRef]
4. Di Franco, C.; Buttazzo, G. Energy-aware coverage path planning of UAVs. In *Proceedings of the IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC), Vila Real, Portugal, 8–10 April 2015*; IEEE: New York, NY, USA, 2015; pp. 111–117. [CrossRef]
5. Brown, S.; Waslander, S.L. The constriction decomposition method for coverage path planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Daejeon, Republic of Korea, 9–14 October 2016*; IEEE: New York, NY, USA, 2016; pp. 3233–3238. [CrossRef]
6. Lien, J.M.; Amato, N.M. Approximate convex decomposition of polygons. *Comput. Geom. Theory Appl.* **2006**, *35*, 100–123. [CrossRef]
7. Liu, Z.; Sengupta, R.; Kurzhanskiy, A. A power consumption model for multi-rotor small unmanned aircraft systems. In *Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS), Miami, FL, USA, 13–16 June 2017*; IEEE: New York, NY, USA, 2017; pp. 310–315. [CrossRef]
8. Jia, B.; Gao, Z.; Jing, J.; Huang, B.; Liu, S.; Muhammad, K.; Rodrigues, J.J.P.C. Coverage path planning for IoUAVs with tiny machine learning in complex areas based on convex decomposition. *IEEE Internet Things J.* **2024**, *11*, 21103–21111. [CrossRef]
9. Datsko, D.; Nekovar, F.; Penicka, R.; Saska, M. Energy-aware multi-UAV coverage mission planning with optimal speed of flight. *IEEE Robot. Autom. Lett.* **2024**, *9*, 2893–2900. [CrossRef]
10. Coombes, M.; Fletcher, T.; Chen, W.H.; Liu, C. Decomposition-based mission planning for fixed-wing UAVs surveying in wind. *J. Field Robot.* **2020**, *37*, 440–465. [CrossRef]
11. Huang, W.H. Optimal line-sweep-based decompositions for coverage algorithms. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Seoul, Republic of Korea, 21–26 May 2001*; IEEE: New York, NY, USA, 2001; pp. 27–32. [CrossRef]
12. Li, Y.; Chen, H.; Er, M.J.; Wang, X. Coverage path planning for UAVs based on enhanced exact cellular decomposition method. *Mechatronics* **2011**, *21*, 876–885. [CrossRef]
13. Bochkarev, S.; Smith, S.L. On minimizing turns in robot coverage path planning. In *Proceedings of the IEEE International Conference on Automation Science and Engineering (CASE), Fort Worth, TX, USA, 21–25 August 2016*; IEEE: New York, NY, USA, 2016; pp. 1237–1242. [CrossRef]
14. Oksanen, T.; Visala, A. Coverage path planning algorithms for agricultural field machines. *J. Field Robot.* **2009**, *26*, 651–668. [CrossRef]
15. Hertel, S.; Mehlhorn, K. Fast triangulation of simple polygons. In *Proceedings of the 4th International Conference on Foundations of Computation Theory, Borgholm, Sweden, 21–27 August 1983*; Springer: Berlin/Heidelberg, Germany, 1983; pp. 207–218. [CrossRef]
16. Torres, M.; Pelta, D.A.; Verdegay, J.L.; Torres, J.C. Coverage path planning with unmanned aerial vehicles for 3D terrain reconstruction. *Expert Syst. Appl.* **2016**, *55*, 441–451. [CrossRef]
17. Nielsen, L.D.; Sung, I.; Nielsen, P. Convex decomposition for a coverage path planning for autonomous vehicles: Interior extension of edges. *Sensors* **2019**, *19*, 4165. [CrossRef]
18. Tang, G.; Tang, C.; Zhou, H.; Claramunt, C.; Men, S. R-DFS: A coverage path planning approach based on region optimal decomposition. *Remote Sens.* **2021**, *13*, 1525. [CrossRef]
19. Kusnur, T.; Likhachev, M. Complete, decomposition-free coverage path planning. In *Proceedings of the IEEE International Conference on Automation Science and Engineering (CASE), Mexico City, Mexico, 20–24 August 2022*; IEEE: New York, NY, USA, 2022; pp. 1431–1437. [CrossRef]
20. Coombes, M.; Chen, W.H.; Liu, C. Boustrophedon coverage path planning for UAV aerial surveys in wind. In *Proceedings of the 2017 International Conference on Unmanned Aircraft Systems (ICUAS), Miami, FL, USA, 13–16 June 2017;* IEEE: New York, NY, USA, 2017; pp. 1563–1571. [CrossRef]
21. Escobar, L.; Pereira, G.A.S. Energy-aware coverage path planner for multirotor UAVs. In *Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS), Charlotte, NC, USA, 14–17 May 2025*; IEEE: New York, NY, USA, 2025; pp. 777–784. [CrossRef]
22. Hong, D.; Lee, S.; Cho, Y.H.; Baek, D.; Kim, J.; Chang, N. Least-energy path planning with building accurate power consumption model of rotary unmanned aerial vehicle. *IEEE Trans. Veh. Technol.* **2020**, *69*, 14803–14817. [CrossRef]
23. Michel, N.; Wei, P.; Kong, Z.; Lin, X. Energy-optimal unmanned aerial vehicles motion planning and control based on integrated system physical dynamics. *ASME J. Dyn. Syst. Meas. Control* **2023**, *145*, 041002. [CrossRef]
24. Theile, M.; Bayerlein, H.; Nai, R.; Gesbert, D.; Caccamo, M. UAV coverage path planning under varying power constraints using deep reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 24 October 2020–24 January 2021*; IEEE: New York, NY, USA, 2020; pp. 1444–1449. [CrossRef]
25. Xiong, Z.; Han, C.; Wang, X.; Gao, L. A coverage path planning method with energy optimization for UAV monitoring tasks. *J. Low Power Electron. Appl.* **2025**, *15*, 39. [CrossRef]

26. Fevgas, G.; Lagkas, T.; Argyriou, V.; Sarigiannidis, P. Coverage path planning methods focusing on energy efficient and cooperative strategies for unmanned aerial vehicles. *Sensors* **2022**, *22*, 1235. [CrossRef]
27. Jensen-Nau, K.R.; Hermans, T.; Leang, K.K. Near-optimal area-coverage path planning of energy-constrained aerial robots with application in autonomous environmental monitoring. *IEEE Trans. Autom. Sci. Eng.* **2021**, *18*, 1453–1468. [CrossRef]
28. Balampanis, F.; Maza, I.; Ollero, A. Spiral-like coverage path planning for multiple heterogeneous UAS operating in coastal regions. In *Proceedings of the International Conference on Unmanned Aircraft Systems (ICUAS), Miami, FL, USA, 13–16 June 2017*; IEEE: New York, NY, USA, 2017; pp. 617–624. [CrossRef]
29. Khaneghaei, M.; Ebrahimi, B.; Asadi, D.; Tutsoy, O.; Nabavi-Chashmi, S.Y.; Haghighi, H. Real-time coverage path planning for fixed-wing aerial robots using partial gradient-based MPC and augmented Dubins trajectories. *Aerospace* **2026**, *13*, 713. [CrossRef]
30. Mier, G.; Valente, J.; de Bruin, S. Fields2Cover: An open-source coverage path planning library for unmanned agricultural vehicles. *IEEE Robot. Autom. Lett.* **2023**, *8*, 2166–2172. [CrossRef]
31. Zeng, Y.; Xu, J.; Zhang, R. Energy Minimization for Wireless Communication with Rotary-Wing UAV. *IEEE Trans. Wirel. Commun.* **2019**, *18*, 2329–2345. [CrossRef]
32. Bähnemann, R.; Lawrance, N.; Chung, J.J.; Pantic, M.; Siegwart, R.; Nieto, J. Revisiting boustrophedon coverage path planning as a generalized traveling salesman problem. In *Field and Service Robotics*; Springer Proceedings in Advanced Robotics; Springer: Singapore, 2021; pp. 277–290. [CrossRef]
33. Di Franco, C.; Buttazzo, G. Coverage path planning for UAVs photogrammetry with energy and resolution constraints. *J. Intell. Robot. Syst.* **2016**, *83*, 445–462. [CrossRef]
34. Mier, G.; Casado Faulí, A.M.; Valente, J.; de Bruin, S. Fields2Benchmark: An open-source benchmark for coverage path planning methods in agriculture. *Smart Agric. Technol.* **2025**, *12*, 101156. [CrossRef]
35. Nesbit, P.R.; Hugenholtz, C.H. Enhancing UAV–SfM 3D model accuracy in high-relief landscapes by incorporating oblique images. *Remote Sens.* **2019**, *11*, 239. [CrossRef]
36. Zhang, H.; Aldana-Jague, E.; Clapuyt, F.; Wilken, F.; Vanacker, V.; Van Oost, K. Evaluating the potential of post-processing kinematic (PPK) georeferencing for UAV-based structure-from-motion (SfM) photogrammetry and surface change detection. *Earth Surf. Dyn.* **2019**, *7*, 807–827. [CrossRef]
37. Han, S.; Han, D. Enhancing direct georeferencing using real-time kinematic UAVs and structure from motion-based photogrammetry for large-scale infrastructure. *Drones* **2024**, *8*, 736. [CrossRef]
38. Jo, I.; Lee, Y.; Ham, N.; Kim, J.; Kim, J.J. A quantitative evaluation of UAV flight parameters for SfM-based 3D reconstruction of buildings. *Appl. Sci.* **2025**, *15*, 7196. [CrossRef]
39. Krupke, D. Near-optimal coverage path planning with turn costs. In *Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX), Alexandria, VA, USA, 7–8 January 2024*; SIAM: Philadelphia, PA, USA, 2024; pp. 118–132. [CrossRef]
40. Georgiou, K.; Kundu, S.; Prałat, P. The Fagnano triangle patrolling problem. *Discret. Math. Theor. Comput. Sci.* **2025**, *27*, 4. [CrossRef]
41. Coumans, E.; Bai, Y. PyBullet, a Python Module for Physics Simulation for Games, Robotics and Machine Learning. 2016–2021. Available online: http://pybullet.org (accessed on 1 October 2026).
42. Panerati, J.; Zheng, H.; Zhou, S.; Xu, J.; Prorok, A.; Schoellig, A.P. Learning to Fly—A Gym Environment with PyBullet Physics for Reinforcement Learning of Multi-agent Quadcopter Control. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Prague, Czech Republic, 27 September–1 October 2021*; IEEE: New York, NY, USA, 2021; pp. 7512–7519. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.