Time-Constrained Coverage Path Planning for UAV Search Applications

Alec Luterman[1], Zachary Bortoff[2], Stephen M. Nogar[3], Derek A. Paley[4]

***Abstract*— Traditional coverage path planning methods for unmanned aerial vehicles (UAVs) take an overly simplistic look at the sensor footprint of the camera, resulting in inefficient path plans that waste significant coverage and are not optimal in travel time. We propose a coverage path planning algorithm based on generating and sequencing a set of stationary vantage points that the UAV will travel to and capture imagery. These vantage points ensure that the coverage path plan achieves a spatial resolution threshold throughout the entire search domain while minimizing the amount of unnecessary excess coverage both inside and outside of the search domain. We also propose a routing method for maximizing the portion of the search domain we can cover when faced with a maximum mission time constraint. Simulation testing shows the improvement of our method in both coverage efficiency and total travel time compared to lawnmower-based coverage path plans. Experimental testing details how coverage path plans based on stationary coverage can degrade when the UAV's onboard camera does not have a level attitude.**

I. Introduction

During mass-casualty events, locating people in need of triage is an important task in order for medics to save as many lives as possible. Unmanned aerial vehicles (UAVs) are a useful tool for this task as they are able to cover regions of interest faster than humans while also being able to inspect regions that humans are unable to, keeping medics out of harm's way [1], [2]. Having UAVs search for casualties autonomously can also free up personnel to start treating casualties instead of trying to locate them. In this work, we focus on a single-agent quadcopter UAV.

The task of generating a path plan for a UAV to cover an entire region of interest is referred to as coverage path planning [3]. This typically consists of finding a path such that the sensor footprint of the UAV covers the entire search domain. The sensor footprint of the UAV at a given pose is the region on the ground that the sensor covers [4]. For a UAV with an attached RGB camera, the sensor footprint at a given pose is calculated by the region formed by the point of intersection on the ground plane from the ray emanating from each pixel in the image using the pinhole camera model. The sensor footprint is directly determined by the camera's intrinsic parameters, such as the focal length and lens distortion, as well as the extrinsic parameters of the camera. Existing coverage path planning methods that account for the sensor footprint do not precisely model the non-uniform distribution of coverage throughout it [5]–[7].

Boustrophedon coverage path plans, commonly referred to as lawnmower patterns, are one of the most common forms of coverage path planning [3], [8]. Unfortunately, they result in uneven distributions of coverage with many regions of the search domain being oversampled while others are undersampled [4]. Motivated by this, we want to generate a coverage path plan that results in an even distribution of coverage. In particular, given a spatial resolution threshold in centimeters per pixel, we want to guarantee that every portion of the search domain achieves this threshold while minimizing the amount of unnecessary additional coverage and minimizing the amount of coverage outside of the search domain. Due to the degradation of object detection quality with a moving camera [9], we also want a coverage path plan that is based on a set of waypoints at which the UAV will be stationary, called vantage points, to perform inspection. Being stationary also minimizes the effect of latency in coordinate frame transformations, thus increasing the localization accuracy of detected objects into the world frame.

Similar to our approach, [10] models coverage path planning as generating and routing a set of vantage points for a UAV to perform inspection at. However, they take a very simplistic view of the sensor footprint, modeling it as a visibility cone. For the task of infrastructure inspection, [11] uses the Generalized Traveling Salesman Problem to ensure complete coverage of a structure of interest by generating vantage points online. [12] defines a distributed control algorithm for coverage from multiple robot platforms based on an objective function that minimizes the area per pixel, a measure of spatial resolution. The major differences between our work and [12] is that our formulation guarantees complete coverage of the search domain in the planned path and our focus is on time-constrained search applications as opposed to persistent monitoring.

The contributions of this paper are (1) a method for coverage path planning that generates a set of vantage points such that their collective coverage results in every portion of the search domain achieving a desired spatial resolution as efficiently as possible; (2) a formulation of an integer linear program that is a variant of the Orienteering Problem to maximize the proportion of the search domain achieving a spatial resolution threshold within a strict time limit; and (3) simulation and experimental results showing the improvement in coverage efficiency of our coverage path

[1]Alec Luterman is with the Department of Computer Science, University of Maryland, College Park, MD, USA. amluter@umd.edu

[2]Zachary Bortoff is with the Department of Aerospace Engineering, University of Maryland, College Park, MD, USA. zbortoff@umd.edu

[3]Stephen M. Nogar is with the DEVCOM Army Research Laboratory stephen.m.nogar.civ@army.mil

[4]Derek A. Paley is with the Department of Aerospace Engineering and the Institute for Systems Research, University of Maryland, College Park, MD, USA. dpaley@umd.edu

This research was supported by Army Cooperative Agreement W911NF2120076

planning algorithm compared to lawnmower patterns and how coverage from stationary vantage points is affected by real-world factors.

Section 2 introduces relevant constraint optimization problems for both the generation and routing of vantage points. Section 3 outlines our approach to coverage path planning. Section 4 describes our experimental setup and provides results from both simulation and experimental testing evaluating the efficiency of coverage from our coverage path planning method. Section 5 concludes the paper.

II. Background

A constraint optimization problem is a mathematical problem in which the goal is to maximize or minimize an objective function of some set of decision variables x, while satisfying a set of equality and inequality constraints on x [13]. The general form of this problem can be seen in Equation (1), where x is the set of decision variables and $f(x)$ is the objective function to be minimized. The equality constraints are the functions $C_i(x), \forall i \in \mathcal{E}$ and the inequality constraints are the functions $C_i(x), \forall i \in \mathcal{I}$, where $\mathcal{E}$ and $\mathcal{I}$ are the sets of constraint indices [13]. The problem is

$$\begin{aligned} \text{Min} \quad & f(x) \\ \text{s.t.} \quad & x \in \mathbb{R}^n \\ & C_i(x) = 0 \quad \forall i \in \mathcal{E} \\ & C_i(x) \geq 0 \quad \forall i \in \mathcal{I} \end{aligned} \tag{1}$$

Of particular relevance to this paper is the subset of constraint optimization problems known as integer linear programming (ILP), where the objective function and constraints are linear functions of the decision variables x, which are constrained to be integer-valued. This paper utilizes an ILP called the Set Cover Problem [14], which involves a universe of elements $U = \{e_i\}_{i=1}^{n}$ and m subsets $s_j \in S$, where $s_j \subseteq U \ \forall j \in \{1, \ldots, m\}$. The goal is to find the minimum number of subsets s_j to cover the entire universe of elements. A formulation of the Set Cover Problem is

$$\begin{aligned} \text{Min} \quad & \sum_{j=1}^{m} x_j \\ \text{s.t.} \quad & \sum_{j:e_i \in s_j} x_j \geq 1 \quad \forall i \in \{1, \ldots, n\} \\ & x_j \in \{0, 1\} \quad \forall j \in \{1, \ldots, m\} \end{aligned} \tag{2}$$

Another ILP we explore is known as the Vehicle Routing Problem (VRP). These are a class of constraint optimization problems with the goal of determining routes for vehicles to visit a set of customers [15]. To model real-world problems, variants of the VRP can be created by adding constraints, such as requiring every single customer to be visited, or using different objective functions such as minimizing distance or maximizing profit. In this paper, we focus on single-agent VRPs.

A well-known variant of the VRP is the Traveling Salesman Problem (TSP), which has the goal of finding a tour of cities that visits each city exactly once while minimizing the total travel cost over all potential routes [16]. Depending on the application, travel cost could be temporal, monetary, distance-based, or some other type of metric.

However, what if it is not in our best interest to travel to every single city? For example, what if the monetary cost to travel to a distant city outweighs the potential benefit of visiting that city? The nominal TSP is not able to handle this variant, which leads to the Prize-Collecting TSP (PC-TSP). In the PC-TSP variant, every city now has a prize for being visited and a penalty if it is not. The goal is to find a route and selection of cities such that the sum of costs and penalties is minimized while achieving a reward threshold W_0 [17].

What if we have limited travel time? How can we maximize our profit? These questions lead to the Orienteering Problem. In this problem, each city has a non-negative reward for visiting and we have a maximum time limit $T_{\max}$. The objective is to find the route and selection of cities to maximize the reward gained [18]–[20]. A close variant of this problem to our work is the Clustered Coverage Orienteering Problem (CCOP) [21], with the primary distinction to our work being the grouping of sampling sites into clusters in the CCOP, with the requirement to visit every cluster and visit a minimum number of sampling sites within each cluster. In addition, their application to water sampling with unmanned surface vehicles requires a capacity budget for the volume of water samples that the robot can carry, which does not have an analog in our domain.

None of these existing VRP variants can be directly used for our task of maximizing the proportion of the search domain we are able to cover within a time limit. The TSP has no mechanism to support a reward-based objective, whereas the reward structure for the PC-TSP and the Orienteering Problem cannot support the reward for a given city being dependent on the other cities we visit. The CCOP has a reward structure that is able to handle the nuances of UAV vantage point coverage, but many of the domain-specific constraints of this formulation do not fit our task. This motivates the formulation of a novel VRP variant that can correctly handle the interdependence of coverage from vantage points, maximizing the area in the search domain we can cover within a time limit.

III. Vantage Point Coverage Path Planning

This section outlines a new coverage path planning algorithm based on generating a set of vantage points that ensure the entire search domain is covered with a guarantee on achieving a spatial resolution threshold throughout every portion of the search domain. Each vantage point is a pose at which the UAV will be stationary for a short duration to perform inspection. Figure 1 shows how the spatial resolution varies throughout the sensor footprint of the camera from an example vantage point. Ideally these vantage points also maximize the efficiency of coverage by minimizing the amount of coverage that is wasted outside of the search domain and by minimizing the portion of the search domain that is covered by multiple vantage points. Given a set of

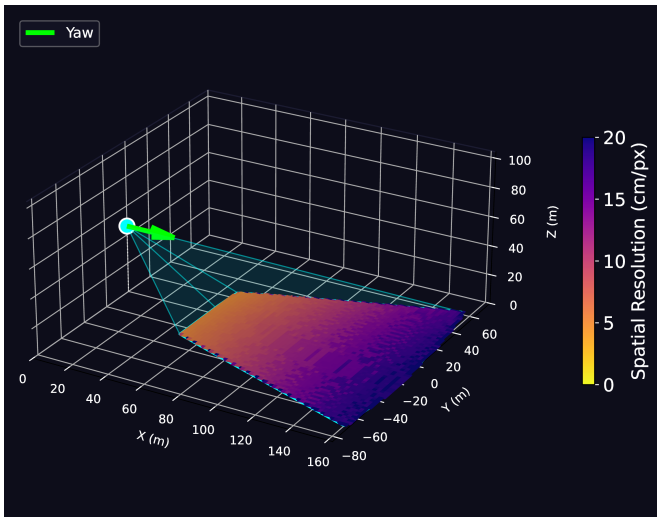

Fig. 1: Example vantage point at an altitude of 46 m with a camera pitched at 60° from nadir to the direction of travel. Lens distortion is accounted for using the Plumb Bob distortion model. The spatial resolution throughout the entire sensor footprint is shown.

vantage points, we determine the coverage path plan by selecting and sequencing a subset of those vantage points.

A. Vantage Point Generation

Given a 2D search domain defined by a set of perimeter vertices, a calibrated camera attached to a UAV, a desired search altitude, and a spatial resolution threshold G in cm/px, we want to generate a set of vantage points such that every $1\,\text{m} \times 1\,\text{m}$ cell in the search domain achieves a minimum number of pixel intersections with the ground plane from the sensor footprint of at least one vantage point. The number of pixel intersections per cell, R, is calculated via $R = (100/G)^2$. Assume that the search domain we are trying to cover is occlusion-free, a known oversimplification for many tasks. Also assume a flat Earth.

We model the task of generating vantage points using the Set Cover Problem. The universe of elements that we are trying to cover is the set P of all cells in the approximate cellular decomposition of the search domain. To construct the family of subsets, we generate a set of candidate vantage points, choosing the regions of the sensor footprint from each that achieve a minimum of R pixel intersections with the ground plane; each subset represents one candidate vantage point. Formally, define $C = \{1, \ldots, m\}$ as the set of indices representing our candidate vantage points where x_j is a binary variable representing if vantage point j is selected. Let $SF(j)$ be the effective sensor footprint for vantage point j, which is the cells in the sensor footprint that achieve a minimum of R pixel intersections with the ground plane. For each cell p in P, define N_p as the set of vantage points that can be chosen to cover cell p, i.e.,

$$N_p = \{j \in C \mid SF(j) \text{ covers cell } p\} \tag{3}$$

We use the ILP formulation in Equation (4) to generate the minimum cardinality set of vantage points such that every cell in the search domain satisfies the minimum pixel intersection threshold, i.e.,

$$\begin{aligned} \text{Min} \quad & \sum_{j=1}^{m} x_j \\ \text{s.t.} \quad & \sum_{j \in N_p} x_j \geq 1 \quad \forall p \in P \\ & x_j \in \{0,1\} \quad \forall j \in C \end{aligned} \tag{4}$$

The candidate vantage points are composed by sampling variations of (x, y) position and yaw within the search domain. Since we hold the flight altitude constant across all vantage points, the only transformations required between the sensor footprints of any two given candidate vantage points are translation and rotation. Note that the effective sensor footprint is invariant under rotation and translation in the XY-plane. This motivates the use of a template effective sensor footprint to greatly reduce the number of computations needed for constructing our subsets, as we only need to calculate the pixel intersections with the ground plane for one vantage point. For any given candidate vantage point, we can calculate its effective sensor footprint by applying a transformation to our template effective sensor footprint.

We use the point $(0, 0, \{\text{search altitude}\})$ with a yaw of $0°$ as the vantage point used to construct our template effective sensor footprint, following the method in [4] to calculate where each pixel intersects the ground plane, before binning these points of intersection into $1\,\text{m} \times 1\,\text{m}$ cells to store pixel counts.

One approach to solving Equation (4) is by using a greedy algorithm [14]. Until every cell in the search domain is covered, we keep selecting the vantage point that will cover the largest number of cells not already covered by the previously selected vantage points. Note that this algorithm provides no guarantee on the cardinality of the generated set of vantage points being optimal. Another choice is using an ILP solving library, such as the CP-SAT solver from Google OR-Tools [22], to solve for the optimal or near-optimal cardinality set of vantage points.

Algorithm 1 below gives an overview of generating a set of vantage points such that every cell in the search domain satisfies the spatial resolution threshold.

Algorithm 1 Vantage Point Generation

1: Approximate cellular decomposition of search domain
2: Generate candidate vantage points
3: Generate template effective sensor footprint
4: Generate effective sensor footprint for each candidate vantage point
5: Solve Set Cover Problem in Equation (4)

B. Vantage Point Routing

Once we generate a set of vantage points, we need to determine the route for our coverage path plan. In a simple

case where we do not have any mission time constraints, using the TSP to sequence all selected vantage points is the best method as this will minimize the total travel time. However, if we have a maximum mission time constraint, such as a limited battery or a regulated flight time restriction, it may not be possible to visit all vantage points. When selecting a subset of vantage points, the amount of additional coverage any given vantage point will add depends on the coverage of all other selected vantage points. In addition, not every vantage point provides the same amount of independent coverage due to some coverage wasted outside of the search domain due to its placement and orientation.

Given a set of vantage points C', a maximum mission time $T_{\max}$, a vantage point inspection duration s, and an estimate for the travel time $t_{i,j}$ between each pair of vantage points (i, j), we formulate the Coverage Orienteering Problem with the goal of maximizing the portion of the search domain that achieves a minimum of R pixel intersections from at least one vantage point within $T_{\max}$. For each cell p in our search domain, let y_p be a binary variable whose value is determined by whether or not our final route will fully cover cell p. z_j is a binary variable representing if vantage point j is selected for our final route. Denote $d \in C'$ as our depot vantage point which we will start our search from and must return to and end the mission before $T_{\max}$. Note that we perform inspection at d when we first travel to it but not on our return. Denote $C^* = C' \setminus \{d\}$ as the set of all non-depot vantage points. Equation (5) below shows the formulation of the Coverage Orienteering Problem.

$$
\begin{aligned}
\text{Max.} \quad & \sum_{p \in P} y_p \\
\text{s.t.} \quad & v_{i,j} \in \{0,1\} \quad \forall i \neq j \in C' \\
& z_j \in \{0,1\} \quad \forall j \in C' \\
& y_p \in \{0,1\} \quad \forall p \in P \\
& y_p \leq \sum_{j \in C' \cap N_p} z_j \quad \forall p \in P \\
& y_p \geq z_j \qquad \forall j \in C' \cap N_p, \forall p \in P \\
& z_d = \sum_{j \in C^*} v_{d,j} = \sum_{i \in C^*} v_{i,d} = 1 \\
& z_k = \sum_{i \in C'} v_{i,k} = \sum_{j \in C'} v_{k,j} \quad \forall k \in C^* \\
& \sum_{i \in C'} \sum_{j \in C'} v_{i,j}(t_{i,j} + s) \leq T_{\max} \\
& u_i - u_j + 1 \leq (|C'| - 1)(1 - v_{i,j}) \\
& \qquad \forall i, j \in C^* \\
& 2 \leq u_j \leq |C'| - 1 \quad \forall j \in C^*
\end{aligned}
\tag{5}
$$

While Equation (5) is generalized to work for any set of vantage points, C' is intended to be the set of vantage points generated by Algorithm 1.

IV. Simulation and Experimental Results

This section evaluates our vantage point coverage path planning algorithm compared to lawnmower-derived path plans with stops for inspection.

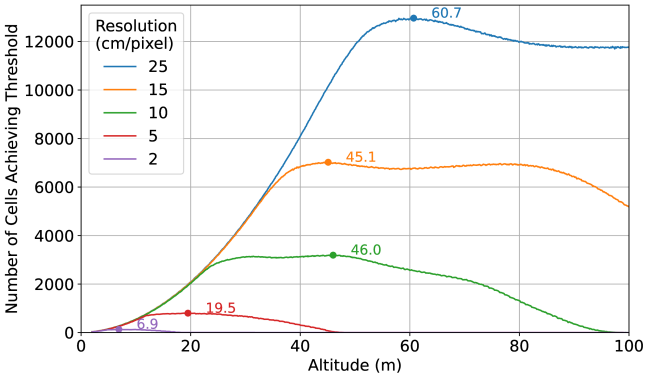

Fig. 2: Number of cells in the sensor footprint of the Grazer-A camera that achieve different spatial resolution thresholds at different flight altitudes. The camera has a pitch of 60° from the ground plane to the UAV's direction of travel.

A. Experimental Setup

To generate flight testing results, we utilize a VOXL2 based UAV paired with a mRo Control Zero H7 flight controller called the ARL Grazer-A [23], [24]. The onboard camera used for high resolution imagery is configured at a pitch of 60° from the ground plane to the direction of travel. It is important to note that this camera is fixed to the UAV and is not on a gimbal. All real-world flight testing was performed at the University of Maryland UAS Research and Operations Center test facility in southern Maryland. For simulation testing, we utilize a digital twin of the Grazer-A and this test facility in a high visual fidelity Unity-based simulation environment [23], [24]. We record the UAV's pose at each vantage point, which we then use to calculate the actual coverage obtained at the vantage point. The vantage point hold duration and transit speed were held constant for all experiments at 5 seconds and 5 m/s respectively.

We use lawnmower-derived path plans with stops for inspection to serve as a reference for our coverage path planning methods. For analysis on vantage point coverage, the track spacing was determined based on the minimum width of the effective sensor footprint and the distance between stops is the height of the effective sensor footprint. This configuration ensures complete coverage and minimizes overlapping coverage. The direction of travel was selected to be parallel to the longest dimension of the search domain. For time-constrained routing analysis, the track spacing was instead chosen to maximize the amount of coverage that we can obtain within the maximum mission time constraint, minimizing overlap.

B. Vantage Point Coverage Analysis

This section evaluates the efficiency of coverage path plans using Algorithm 1 to generate vantage points. Testing was performed in a high-fidelity simulation environment. For solving Equation (4), we compare using a greedy algorithm versus using the CP-SAT solver from Google OR-Tools

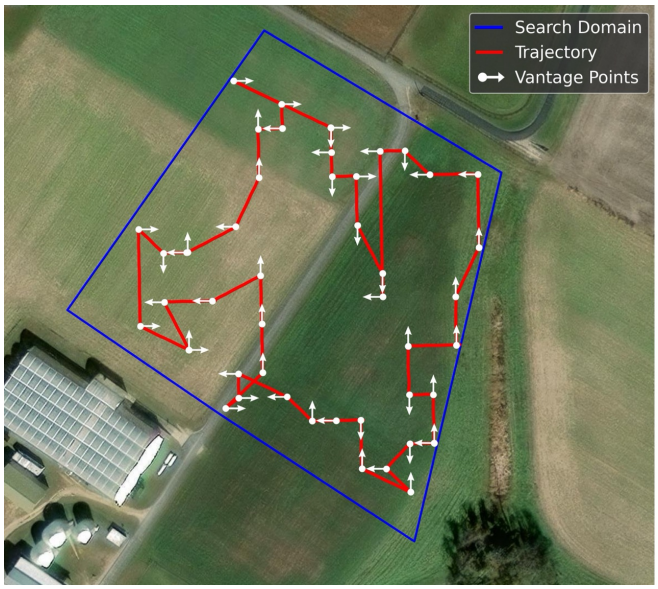

Fig. 3: Vantage point coverage path plan flown during simulation testing for a spatial resolution threshold of 5 cm/px at an altitude of 19.5 m.

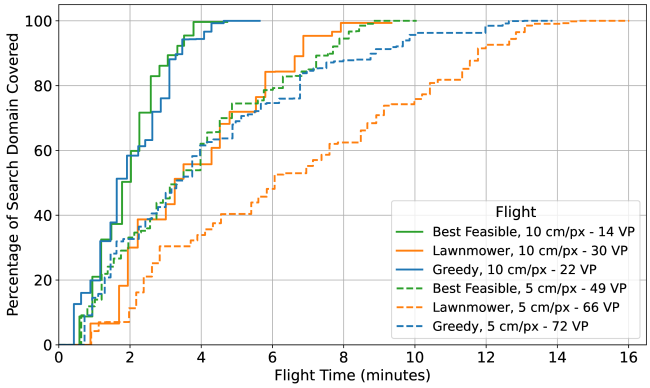

Fig. 4: Percentage of the search domain covered over time for different coverage path planning methods and spatial resolution thresholds in a simulation environment. The number next to each flight is the number of vantage points.

[22], taking the best feasible solution after two minutes of solving time to mimic a real-world operation. Note that we consistently saw an optimality gap of less than three with this configuration for CP-SAT. For both solving methods, vantage points are sequenced by solving the TSP. Candidate vantage points were sampled every 10 meters in both directions of the XY-plane inside of the search domain, with yaws of 0°, 90°, 180°, and 270°. We also compare with a lawnmower-derived baseline as defined earlier. We evaluate these three methods at spatial resolution thresholds of 5 cm/px and 10 cm/px, with the search altitude chosen to maximize the size of the effective sensor footprint based on Figure 2. This Figure shows the number of the cells in the effective sensor footprint of the Grazer-A camera at different altitudes and spatial resolution thresholds. For a given spatial resolution threshold, we choose the altitude resulting in the largest number of cells achieving that threshold. Figure 3 shows the search domain used for this set of testing and an example vantage point coverage path plan.

Based on data collected from simulation, Figure 4 shows that the vantage point coverage path plans achieve full coverage of the search domain quicker compared to the lawnmower-based path plans. At a spatial resolution threshold of 5 cm/px, with a corresponding altitude of 19.5 m, the path plan generated using the best feasible solution took 10 minutes and 3 seconds, the greedy path plan took 13 minutes and 52 seconds, and the lawnmower path plan took 16 minutes and 1 second. These results also extend to the set of flights at a threshold of 10 cm/px and an altitude of 46 m, with the path plan from the best feasible solution taking 4 minutes and 44 seconds, the greedy path plan taking 5 minutes and 38 seconds, and the lawnmower path plan taking 9 minutes and 20 seconds.

For the ordering of the two vantage point solving methods, the best feasible solution results in a better solution to solving Equation (4) compared to a greedy solution based on minimizing the number of vantage points. As a result, the greedy path plan has more vantage points to travel to and perform inspection at, increasing the total travel time. When comparing to the lawnmower path plans, even though at 5 cm/px the greedy path plan actually had more vantage points than the lawnmower path plan, many of these vantage points actually had the same (x, y) position, with both path plans having 66 unique vantage point positions. This highlights the inefficiency of the lawnmower pattern.

Not surprisingly, a denser spatial resolution requires more flight time to complete full coverage of the search domain as across all methods, the 5 cm/px flights took longer than their 10 cm/px counterparts. This is because the effective sensor footprint is larger at the optimal altitude for 10 cm/px, thus fewer vantage points are required to cover the entire search domain compared to a spatial resolution of 5 cm/px.

One important thing to note is that coverage is not just a binary metric, but actually has three different categories: unique coverage, overlapped coverage, and overspray coverage. Unique coverage is the regions in the search domain that are covered by exactly one vantage point, sufficiently meeting our spatial resolution threshold. Overlapped coverage describes regions in the search domain that are covered by more than one vantage point. Overspray coverage denotes the regions outside of the search domain that were unintentionally covered.

Figure 5 shows that the method of vantage point selection has a big impact on the portion of the cells seen at vantage points which are classified as unique or overlapped. At both spatial resolution thresholds, the path plan from the best feasible solution to Equation (4) results in the largest percentage of unique coverage compared to the greedy path plan, showing the validity of our objective function. In addition, the usage of Equation (4) to generate vantage points also shows its benefits for minimizing overspray based on the

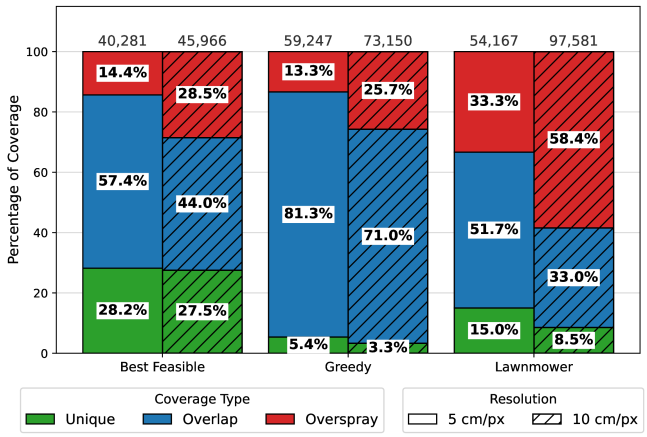

Fig. 5: Characterization of coverage measured from all vantage points from all simulated test flights. The number above each bar is the number of cell instances covered by vantage points, i.e., a cell covered by n vantage points independently counts as n instances.

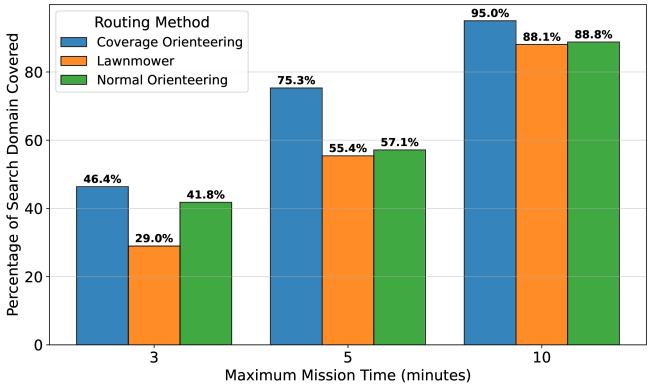

Fig. 6: Percentage of the search domain covered at a spatial resolution of 5 cm/px for different routing methods for different maximum mission times.

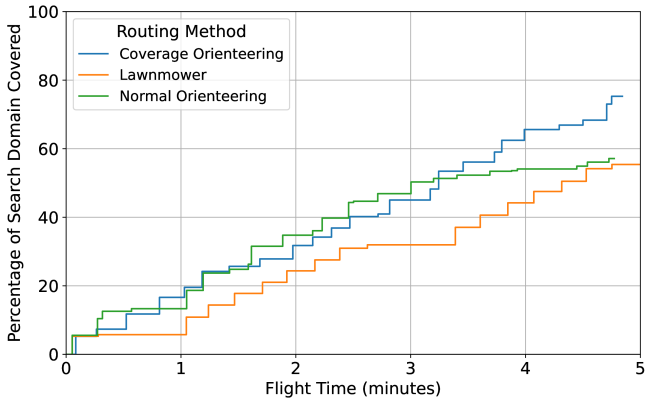

Fig. 7: Percentage of the search domain covered over time for different routing methods for a maximum search time of 5 minutes at a spatial resolution of 5 cm/px.

fact that at both spatial resolution thresholds the lawnmower pattern resulted in significantly more overspray compared to the vantage point-based path plans. One other item of note is that even though for each method the 5 cm/px flight required more than double the number of vantage points, these flights had fewer total cell instances seen. This is due to the size of the effective sensor footprint being significantly smaller at this denser spatial resolution.

C. Vantage Point Routing Analysis

This section compares different routing methods on their ability to maximize the portion of the search domain achieving a spatial resolution threshold within a limited mission time. We evaluate lawnmower coverage path plans, the Orienteering Problem with the objective to maximize the number of vantage points visited (we refer to this as the normal Orienteering Problem), and the Coverage Orienteering Problem at a spatial resolution of 5 cm/px in a search domain of over 9 acres for maximum mission times of 3 minutes, 5 minutes, and 10 minutes. Both variants of the Orienteering Problem were solved using CP-SAT with a maximum solving time of two minutes and taking the best feasible solution. The set of vantage points used was the best feasible solution after two minutes of solving Equation (4). These flights were performed in a simulation environment.

Figure 6 shows the effect that the choice of routing method has on the total portion of the search domain we are able to cover at the desired spatial resolution within the time limit. For all three maximum mission times, the Coverage Orienteering Problem results in the greatest portion of the search domain covered. While not as significant of a difference, the normal Orienteering Problem does result in a larger portion of the search domain being covered compared to the lawnmower path plan. From Figure 7, which had a time limit of 5 minutes, we can see that the Coverage Orienteering Problem results in the most consistent gain in coverage, as the normal Orienteering Problem's gain in coverage decreases over time, and the lawnmower has long stretches of no gain in coverage due to travel between tracks, highlighting the benefit of the Coverage Orienteering Problem.

The reason for the difference between these Orienteering Problem variants is due to the types of vantage points they are incentivized to select. From Figure 8, which had a time limit of 10 minutes, we can see that the Coverage Orienteering Problem chooses vantage points more spread out across the search domain and also chooses fewer vantage points on the boundary of the search domain while the normal Orienteering Problem is incentivized to select vantage points closer together resulting in high amounts of overlap. The normal Orienteering Problem also results in the highest number of vantage points selected as with the 10 minute time limit, the Coverage Orienteering Problem selected 51 vantage points while the normal Orienteering Problem selected 55. At the lower time limits, this trend also holds. These results show that optimizing for the number of vantage points within a time limit results in a smaller portion of the search domain covered compared to optimizing directly for the area of

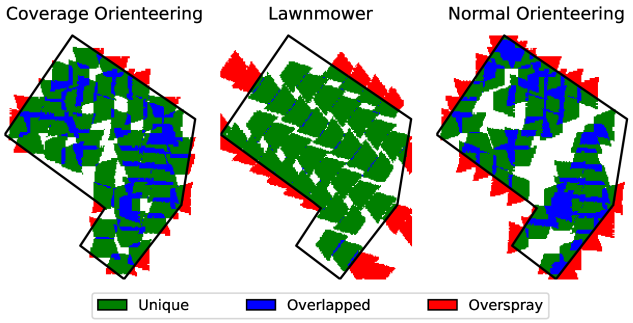

Fig. 8: Map of the characterization of coverage for different time-limited routing methods for a maximum mission time of 10 minutes at a spatial resolution of 5 cm/px.

the search domain covered using the Coverage Orienteering Problem.

D. Experimental Results

For real-world testing, we flew six flights at a spatial resolution threshold of 10 cm/px, varying two factors: the vantage point solving method (greedy versus the best feasible solution after two minutes of solving) and the flight altitude (15 m, 30 m, and 45 m). These flights used the same search domain as our vantage point coverage analysis simulation testing, shown in Figure 3, and were performed over two days in early February 2026, facing winds of over 13 knots at a direction of 350°. Figure 9 shows coverage statistics from these flights.

The most apparent item is that these flights did not achieve full coverage of the search domain, with the greedy coverage path plan achieving a larger portion of coverage compared to the flights from the best feasible path plan for each altitude. At each vantage point, the flight behavior of the Grazer-A prioritizes holding (x, y, z) position over maintaining a level attitude. As a result, when faced with heavy wind, the flight controller will command heavy roll and pitch to hold position, causing a deviation from the orientation of the planned pose of the vantage point. Since the Grazer-A does not have a gimbal, this change in attitude directly affects the attitude of the camera and thus the sensor footprint.

Figure 10 shows how this negatively impacts the total coverage obtained. We can see that the size and shape of the effective sensor footprints from the actual coverage at each vantage point are notably deformed compared to the planned coverage. If we assume level attitude for the camera for the calculation of coverage at each vantage point, all six of these flights actually achieve more than 99.5% coverage of the search domain, showing that the camera's attitude is the primary contributor to the discrepancy of coverage. This is also backed up visually by the minimal gaps in coverage when assuming a level attitude in Figure 10.

This discrepancy between the planned coverage and the real-world coverage is a result of not having a method for the camera to maintain level attitude in flight and would affect any path plan built on stationary points for inspection. This

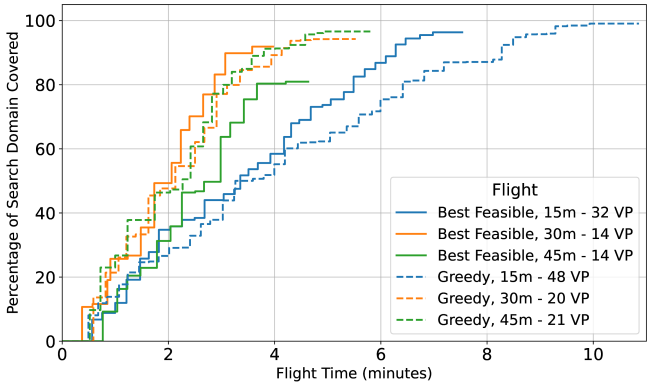

(a) Percentage of the search domain covered over time. The number next to each flight is the number of vantage points.

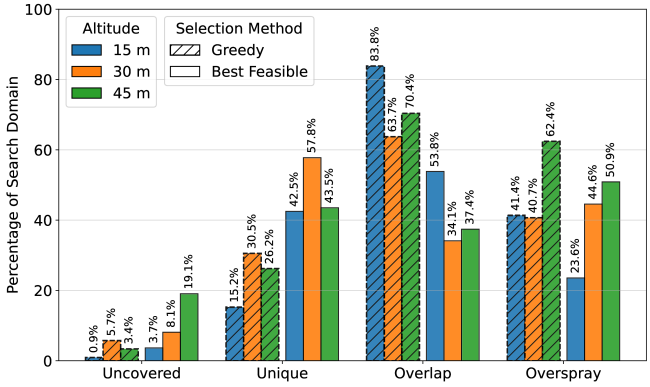

(b) Characterization of coverage relative to the size of the search domain.

Fig. 9: Portion of the search domain covered from all six experimental test flights.

does not detract from the value of using stationary vantage points for coverage path planning given its benefits for target detection and localization, it just means that we need a way to ensure the camera can obtain the planned orientation at each vantage point. A potential fix to this would be using a UAV with a multi-axis gimbal.

V. Conclusion

We propose a new method for UAV coverage path planning that ensures the entire search domain achieves a spatial resolution threshold from the cumulative coverage of a set of vantage points, while minimizing unnecessary additional coverage. When faced with a maximum mission time, we also propose a method for maximizing the portion of the search domain covered using the formulation of a new variant of the Orienteering Problem, called the Coverage Orienteering Problem. In simulated testing we show that our coverage path planning algorithm provides a more efficient distribution of coverage throughout the search domain, with more unique coverage and less wasted coverage compared to lawnmower-based patterns, while also covering a search domain faster. However, in real-world usage, we show how coverage path plans can degrade with deviations from level attitude of the camera.

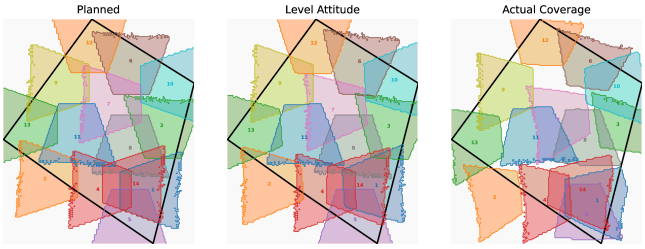

(a) Best feasible solution at 30 m altitude.

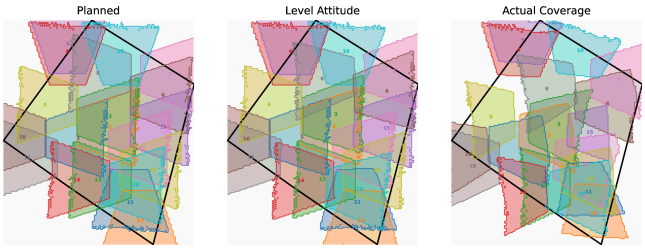

(b) Greedy at 30 m altitude.

Fig. 10: Planned coverage at vantage points versus actual coverage at vantage points assuming level attitude of the camera versus actual coverage at vantage points. Each trapezoid is the effective sensor footprint at the corresponding vantage point, using a spatial resolution threshold of 10 cm/px. These results are from experimental testing.

In ongoing and future work, we are exploring the extension of our coverage path planner to a UAV with a multi-axis gimbal to maintain level attitude for the camera during a search mission. We are also exploring the expansion of dynamic planning to our vantage point coverage path planning method. Based on information gained during search such as the real-time coverage that we obtained, new regions of interest to cover, or existing regions of interest to increase coverage on, we want to update the current coverage path plan we are flying to fulfill these mission goals.

Acknowledgment

We thank Chris Titus, Alex Vrbka, Josh Gaus, Grant Williams, and Rob Neuner from the University of Maryland UAS Research and Operations Center for their help conducting flight experiments. We also thank Dr. Bruce Golden and Dr. Deep Ray for their consultations.

References

[1] M. Lyu, Y. Zhao, C. Huang, and H. Huang, "Unmanned aerial vehicles for search and rescue: A survey," *Remote Sensing*, vol. 15, no. 13, 2023. [Online]. Available: https://www.mdpi.com/2072-4292/15/13/3266

[2] S. M. S. Mohd Daud, M. Y. P. Mohd Yusof, C. C. Heo, L. S. Khoo, M. K. Chainchel Singh, M. S. Mahmood, and H. Nawawi, "Applications of drone in disaster management: A scoping review," *Science Justice*, vol. 62, no. 1, pp. 30–42, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1355030621001477

[3] A. Otto, N. Agatz, J. Campbell, B. Golden, and E. Pesch, "Optimization approaches for civil applications of unmanned aerial vehicles (UAVs) or aerial drones: A survey," *Networks*, vol. 72, no. 4, pp. 411–458, 2018. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.21818

[4] Z. Bortoff, A. Luterman, D. A. Paley, and S. M. Nogar, "Using target detection probability to evaluate area coverage by a UAV," in *2025 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2025, pp. 572–578.

[5] C. Di Franco and G. Buttazzo, "Energy-aware coverage path planning of UAVs," in *2015 IEEE international conference on autonomous robot systems and competitions*. IEEE, 2015, pp. 111–117.

[6] D. Kingston, S. Rasmussen, and L. Humphrey, "Automated UAV tasks for search and surveillance," in *2016 IEEE Conference on Control Applications (CCA)*, 2016, pp. 1–8.

[7] A. Majeed and S. Lee, "A new coverage flight path planning algorithm based on footprint sweep fitting for unmanned aerial vehicle navigation in urban environments," *Applied Sciences*, vol. 9, no. 7, 2019. [Online]. Available: https://www.mdpi.com/2076-3417/9/7/1470

[8] T. M. Cabreira, L. B. Brisolara, and P. R. Ferreira Jr., "Survey on coverage path planning with unmanned aerial vehicles," *Drones*, vol. 3, no. 1, 2019. [Online]. Available: https://www.mdpi.com/2504-446X/3/1/4

[9] M. Yazdi and T. Bouwmans, "New trends on moving object detection in video images captured by a moving camera: A survey," *Computer Science Review*, vol. 28, pp. 157–177, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1574013716301794

[10] J. Zhang and H. Huang, "Occlusion-aware UAV path planning for reconnaissance and surveillance," *Drones*, vol. 5, no. 3, 2021. [Online]. Available: https://www.mdpi.com/2504-446X/5/3/98

[11] H. Dhami, K. Yu, T. Williams, V. Vajipey, and P. Tokekar, "GATSBI: An online GTSP-based algorithm for targeted surface bridge inspection," in *2023 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2023, pp. 1199–1206.

[12] M. Schwager, B. J. Julian, M. Angermann, and D. Rus, "Eyes in the sky: Decentralized control for the deployment of robotic camera networks," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1541–1561, 2011.

[13] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer New York, NY, 2006, ch. 1.

[14] D. P. Williamson and D. B. Shmoys, *The design of approximation algorithms*. Cambridge university press, 2011.

[15] K. Braekers, K. Ramaekers, and I. Van Nieuwenhuyse, "The vehicle routing problem: State of the art classification and review," *Computers Industrial Engineering*, vol. 99, pp. 300–313, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360835215004775

[16] N. Christofides, "Technical note—bounds for the travelling-salesman problem," *Operations Research*, vol. 20, no. 5, pp. 1044–1056, 1972. [Online]. Available: https://doi.org/10.1287/opre.20.5.1044

[17] D. Bienstock, M. X. Goemans, D. Simchi-Levi, and D. Williamson, "A note on the prize collecting traveling salesman problem," *Mathematical programming*, vol. 59, no. 1, pp. 413–420, 1993.

[18] B. L. Golden, L. Levy, and R. Vohra, "The orienteering problem," *Naval Research Logistics (NRL)*, vol. 34, no. 3, pp. 307–318, 1987.

[19] T. Tsiligirides, "Heuristic methods applied to orienteering," *The Journal of the Operational Research Society*, vol. 35, no. 9, pp. 797–809, 1984. [Online]. Available: http://www.jstor.org/stable/2582629

[20] P. Vansteenwegen, W. Souffriau, and D. V. Oudheusden, "The orienteering problem: A survey," *European Journal of Operational Research*, vol. 209, no. 1, pp. 1–10, 2011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221710002973

[21] W. Zhang, K. Wang, S. Wang, and G. Laporte, "Clustered coverage orienteering problem of unmanned surface vehicles for water sampling," *Naval Research Logistics (NRL)*, vol. 67, no. 5, pp. 353–367, 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.21906

[22] L. Perron and F. Didier. (2025) CP-SAT solver. Version 9.12. [Online]. Available: https://developers.google.com/optimization/cp/cp_solver/

[23] W. Cui, A. Shastry, D. A. Paley, and S. Nogar, "Autonomous aerial search and revisit behavior for communication limited environments," in *AIAA SCITECH 2024 Forum*, 2024, p. 1170.

[24] W. Cui, S. S. Poojari, S. S. Abdi, S. M. Nogar, and D. A. Paley, "Long-range autonomy for collaborative aerial search and target revisit," *Journal of Guidance, Control, and Dynamics*, vol. 48, no. 11, pp. 2588–2598, 2025. [Online]. Available: https://doi.org/10.2514/1.G008897