Received XX Month, XXXX; revised XX Month, XXXX; accepted XX Month, XXXX; Date of publication XX Month, XXXX; date of current version XX Month, XXXX.

Digital Object Identifier 10.1109/XXXX.2022.1234567

Terrain-Aware Probabilistic Search Planning for Unmanned Aerial Vehicles

Nathan L. Schomer[1], Member, IEEE, Julie A. Adams[2], Member, IEEE

[1]Collaborative Robotics and Intelligent Systems Institute, Oregon State University, Corvallis, OR 97330 USA
[2]Collaborative Robotics and Intelligent Systems Institute, Oregon State University, Corvallis, OR 97330 USA

Corresponding author: Nathan L. Schomer (email: schomern@oregonstate.edu).

ABSTRACT Mountain search and rescue is a form of emergency response to assist people in austere environments (e.g., extreme terrain, poor weather). Volunteer mountain search and rescue teams in the United States have begun adopting consumer-grade unmanned aerial vehicles to assist a variety of tasks (e.g., search, resource delivery); however, these tools lack the autonomy necessary for the mountain search and rescue teams to fully realize their potential for wide area, aerial search. The unique and tight constraints of mountain search and rescue (e.g., in situ computation, sensor limitations) greatly limit the applicability of recent robotics research. A two-step coverage path planning algorithm that leverages existing viewpoint and path planning approaches was developed to meet the unique needs of mountain search and rescue. Viewpoints were sampled to meet a minimum coverage ratio and assigned priority from a search probability map. The path planning problem was formulated as a clustered traveling salesperson problem, which is solved with a metaheuristic iterative solver. Simulation results inform parameter selection for a series of field experiments. The field experiments demonstrate how the new algorithm can provide resilience against the many compounding factors that make UAV-based mountain search and rescue challenging.

INDEX TERMS Autonomous Aerial Vehicles, Autonomous Robots, Search and Rescue, Disaster Response.

I. INTRODUCTION

Mountain search and rescue (MSAR) in the United States of America (USA) is a volunteer-based emergency response system created to perform search and technical rescue of people in alpine environments. MSAR is distinct from other forms of search and rescue, because it is used when the emergency response requires technical search and rescue, due to terrain or environmental hazards (e.g., steep rock, snow, ice, poor weather). These environments are often remote without vehicular access, which necessitates MSAR personnel carry all required equipment by foot. This equipment includes the rescuer's personal equipment (e.g., food, water, bivouac sack, sleeping bag) as well as any gear necessary to the response (e.g., medical gear, rescue litter, rope rescue gear). This equipment can quickly consume all of rescuer's carrying capacity. MSAR personnel must be discerning when selecting gear to carry into the field. Any new MSAR-related technology must be developed with these size and weight constraints in mind.

MSAR is a subset of wilderness search and rescue. Wilderness search and rescue assists people in remote areas (e.g., forests, mountains, deserts) (Goodrich et al., 2008). This manuscript focuses on wilderness search and rescue operations specifically in mountains (i.e., MSAR). There are many commonalities between different wilderness search and rescue domains; however, MSAR's unique constraints (e.g., vertical terrain, cold temperatures, numerous visual occlusions, high winds) make this distinction necessary.

A 2022 survey of USA MSAR teams about UAV usage and affordability found 83.3% of respondents reported using UAVs (Schomer and Adams, 2022). The respondents reported use of multi-rotor UAVs, but no fixed-wing UAVs. The results show that visible-spectrum cameras are used most frequently, followed closely by thermal cameras. 63.6% of the respondents had less than $10,000 in average annual surplus, while 45.5% had less than $5,000. This limiting funding constrains MSAR teams to low-cost off-the-shelf UAVs. Purchasing a new, $20,000 mapping UAV with LiDAR is not a feasible solution. Additionally, MSAR UAV operations are much riskier in alpine environments. The hardware must be affordable enough that it can be insured/replaced in the event of vehicle loss to be a feasible solution for MSAR, due to these funding constraints.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

©2026 IEEE

A major limitation of UAVs is their limited flight times (Vincent-Lambert et al., 2023). Many multi-rotor UAVs have flights times of 30-40 minutes under ideal environmental conditions (e.g., low wind, warm, low elevation); however, MSAR operations often occur in adverse environmental conditions (e.g., high winds, extreme cold, high elevation), which further reduces reliable battery life. UAVs may only have 15-20 minutes of flight time before having to return to the designated landing area in adverse conditions. MSAR personnel can only carry a limited number of UAV batteries and have no practical means of recharging batteries in the field. UAVs used for MSAR may be unable to complete sensor coverage of the entire search area because of these energy constraints. UAV energy models are too inaccurate to effectively determine path feasibility under real-world conditions (Schomer and Adams, 2024). These limitations require that MSAR personnel have a means of prioritizing high-probability regions early in the search.

Search maps used by MSAR teams do not encode probability, rather they show search priority (Schomer and Adams, 2024). There exist a wealth of solutions to priority-based search problems; however, none currently meet MSAR's unique requirements. Any MSAR-relevant solution must provide complete coverage of the search region, while minimizing energy usage and prioritizing high-probability search regions earlier in the path. This prioritization is necessary to maximize probability of detection if a path is not completed due to insufficient energy; however, strict priority following may lead to highly inefficient (i.e., long) paths for some search maps (e.g., two disjoint high probability regions). Therefore, a method of scaling the trade-off between search area prioritization and energy efficiency is necessary.

This manuscript presents a new approach to MSAR UAV aerial search. Two primary contributions are discussed. A primary contribution is a more computationally efficient viewpoint planning algorithm. Another primary contribution is a path planning algorithm based on a traveling salesperson problem formulation that prioritizes viewpoint clusters based on the rescuer's believed search area probability distribution. Field experiments demonstrate efficacy of these approaches for in situ MSAR coverage path planning. Practical recommendations are provided to support the near-term adoption of UAV autonomy within the MSAR community.

Section 2 provides a review of path planning literature. Section 3 presents a coverage path planning algorithm incorporating rescuer's believed probability distributions. Section 4 discusses experimentation and Section 5 provides results of simulation and field experiments.

II. BACKGROUND

Path planning algorithms calculate a path to move a robot from a start pose to a goal and often begins by decomposing the robot's workspace (i.e., the set of all poses a robot can reach). Workspace dimensions commonly includes the UAV's body pose (e.g., GPS position, altitude, heading) and may include additional degrees of freedom from auxiliary hardware (e.g., camera gimbal). This discreteziation simplifies a continuous problem into a discrete one. A path can be planned over this workspace, optimized for certain metrics or other parameters (e.g., distance, time, minimum obstacle standoff). The resultant path is an ordered set of poses (i.e., waypoints) required to move a robot from a start pose to a goal pose. Coverage Path Planning (CPP) seeks to plan a path allowing the robot to survey an area of interest, achieving a specified level of thoroughness (e.g., capturing at least 95% of the area in camera imagery).

UAV coverage paths over flat terrain with obstacles can be planned using two dimensions (2D). This 2D path planning problem can be solved exactly using geometric methods, such as Boustrophedon cellular decomposition (Choset and Pignon, 1998). However, MSAR operations rarely occur on simple, flat terrain. Three-dimensional (3D) environmental features within the search area must be searched (e.g., mountains, buttes, canyons). The simplest method of incorporating elevation into path planning is via 2.5D path planning (Nam et al., 2017). This method first plans paths over a 2D map, then adds elevation to the planned waypoints by sampling a digital elevation map (DEM). Elevation information greatly expands this method's application space over 2D planning, without adding significant computational cost.

2.5D approaches often assume fixed sensor poses relative to the UAV (e.g., a gimbaled camera is kept in the downward position). This assumption only holds when the terrain is simple and the entire area of interest can be effectively captured by nadir imagery. The viewing angle from a nadir camera will be too acute to effectively sense the region if surfaces in the area are too vertical. Coverage path planning in cases with verticality requires 3D path planning. MSAR typically operates in complex terrain, requiring 3D planning.

3D coverage path planning is useful for many applications where a 3D area must be completely captured by sensor data (e.g., industrial inspection, remote sensing, precision agriculture) (Whitehead and Hugenholtz, 2014) (Muchiri and Kimathi, 2022). There exists both prior research and commercial software to perform 3D inspection tasks. Available software and standard workflows used for these applications are not practical for MSAR applications due to many design features intended for non-remote applications (e.g., complex workflow, required pre-flights, battery charging) (Skydio, 2023) (Bolourian and Hammad, 2020) (Jordan et al., 2018). Additionally, the inherent energy constraints of UAV-based MSAR search preclude the use of 3D coverage algorithms that do not prioritize regions (e.g., commercially available coverage/inspection algorithms, standard geographic survey techniques). Unlimited UAV flight time is not a luxury afforded to MSAR operators (Vincent-Lambert et al., 2023).

Path efficiency and accuracy are critical coverage path planning metrics. Path efficiency includes metrics that dictate the total time or energy a UAV has to expend to complete the path (Tan et al., 2021). Other path efficiency metrics

include 3D path length and total energy consumption. Path accuracy describes how well the calculated coverage path meets the desired path parameters (e.g., total coverage percentage, sensor overlap) (Tan et al., 2021). Coverage path planning algorithms for MSAR aerial search applications must prioritize total energy cost (i.e., path efficiency) due to the limited batteries that rescuers can carry. Computational complexity and convergence time will be limiting factors if in situ planning is required (i.e., the path is planned on a portable device) and planning takes a burdensome amount of time. Interviewed MSAR subject matter experts believe that path planning times of less than 30 minutes will not burden MSAR operations. There are many MSAR tasks that can be performed during path planning (Deboodt, 2023).

A survey established ten categories of coverage path planning algorithms: random walk, chaotic coverage path planner, spanning tree coverage, artificial potential field, sampling-based planning, dynamic programming, greedy graph search, evolutionary, human-inspired, and classical-heuristic algorithms (Tan et al., 2021). Sampling-based and human-inspired coverage path planning algorithms are the algorithm categories best suited for large-scale structural coverage (e.g., building inspection, MSAR aerial search); however, human-inspired algorithms are intended for exploration tasks with limited a priori environmental information (Tan et al., 2021). Sampling-based approaches are useful for MSAR because they scale to large, complex planning tasks. The other eight approaches do not apply to MSAR because they plan online, are unable to handle complex 3D environments, or are too computationally expensive.

The primary goal of coverage path planning for MSAR search applications is capturing the full search area with a camera. Viewpoint-based planning is a sampling-based coverage path planning algorithm that calculates camera viewpoints to achieve total coverage, which can be treated as a set coverage problem, rather than robot poses. The resulting traveling salesperson problem is solved to calculate the most efficient path. This approach decouples the problem into two distinct steps: viewpoint and path planning.

A viewpoint is defined as a single camera pose and orientation. The terrain coverage (i.e., the visible, non-occluded terrain captured by the camera) of a given viewpoint is determined by the camera's extrinsic and intrinsic orientations. A common approach to viewpoint planning is the next best view algorithm, a greedy approach that calculates successive viewpoints using only the robot's current pose and information collected from the robot's sensors (Daudelin and Campbell, 2017). Next best view excels when the robot is tasked with coverage of an unknown structure (i.e., model-free planning); however, this is not the case for MSAR.

MSAR search planning is considered model-based, because of the availability of geographic data (e.g., digital elevation data, topographic maps). MSAR coverage path planning problems can be treated as set coverage problems, rather than exploration tasks, because of this available data. The quality and reliability of DEMs for search planning is highly variable. The US Geographic Survey 3D Elevation Program (3DEP) provides a $\frac{1}{3}$ arc second per pixel (approximately 10 meters per pixel) DEM for the contiguous USA (Stoker and Miller, 2022); however, this map is low-resolution relative to the size of MSAR search subjects and clues, thus it will not capture all relevant terrain occlusions (Polidori and El Hage, 2020). Additionally, MSAR operations occur in terrain that undergoes both seasonal and permanent change. Alpine environments can be covered by meters-deep snow, depending on the season. Additionally, these environments are subject to rapid terrain morphology (e.g., rockfall, wildfires, vegetation growth, glacial melt) (Polidori and El Hage, 2020). Algorithms developed for MSAR search must account for higher model uncertainty than other applications. A possible solution is the long-timescale explore-then-exploit approach that initially captures a model of the area of interest using exploration techniques (Maboudi et al., 2023). This model is used to plan more comprehensive coverage viewpoints of the area of interest (Maboudi et al., 2023). A major limitation of this approach is the increased energy consumption required for the two-step process. MSAR teams often respond to the same locations multiple times. Trained MSAR personnel may be able to capture terrain features to create accurate, up-to-date DEMs.

Model-based viewpoint coverage path planning can be formulated in two steps: viewpoint and path planning. The viewpoint planning step establishes a minimum set of camera viewpoints to cover an area given certain constraints (e.g., minimum distance from terrain, maximum distance from terrain, minimum viewing angle). The path planning step calculates a path between the selected viewpoints. The traveling salesperson problem (TSP) and orienteering problem are two optimization problems commonly used to formulate planning problems when a single agent must visit all, or a subset of a set of points. TSP seeks a minimum cost path between all points (Matai et al., 2010a). The orienteering problem considers a scenario with a set of discrete points of interest, each with an associated reward value. The orienteering problem objective is to maximize the total collected reward, given a limited travel budget. The cost of traveling between points of interest is often proportional to the distance between these points (Vansteenwegen et al., 2011). This problem formulation differs from the MSAR scenario, as it considers a predefined travel budget. MSAR search seeks to obtain complete, or maximum coverage of a search area, given limited, but unknown energy capacity (i.e., UAV flight time). Rescuers will know how many batteries are available for the search; however, there is no effective means of accurately calculating the available flight times. These constraints dictate that TSP is a more appropriate problem formulation for MSAR path planning: a minimum-cost path between all viewpoints must be calculated to maximize the search coverage. Unknown available flight time additionally

negates use of other TSP formulations that calculate subsets of the viewpoints based on a budget (i.e., budgeted TSP).

A coverage path planning algorithm using binary voxel dilation for viewpoint sampling and greedy search was proposed to solve the set cover problem (Jing et al., 2016). The process begins by constructing a voxel model of the area of interest. This model is dilated by the minimum and maximum standoff distances. The minimum dilation is subtracted from the maximum dilation to yield the viewpoint sample space. The maximum standoff distance is constrained by the sensor resolution and desired minimum ground sample distance, where the minimum standoff distance is a safety parameter to ensure minimum distance from terrain collisions. The viewpoint sample space is randomly sampled. An iterative sampling processes is employed until complete coverage of the area of interest is achieved by the sampled viewpoints. These viewpoints must be down-sampled to a minimum coverage set. Jing et. al. treated viewpoint planning as a set coverage problem (SCP) and solved it using a greedy search algorithm. Decoupling viewpoint planning and path planning is a natural choice for MSAR where the search area to be covered does not change significantly between flights, but the viewpoint order may change (depending on which viewpoints are prioritized in a given search scenario).

The terms *waypoints* and *viewpoints* will be used interchangeably; however, the term *vertices* is commonly used in TSP-related literature. There exist many different TSP formulations (Matai et al., 2010b) (Santini et al., 2022) (Manerba et al., 2017). Goodrich et. al. applied TSPs to wilderness search and rescue path planning problems to maximize coverage of a probability distribution in a set amount of time (Goodrich and Lin, 2009).

The clustered traveling salesperson problem (CTSP) divides the waypoints into clusters and requires that all waypoints in each cluster must be visited consecutively (Chisman, 1975). CTSP is analogous to TSP when an arbitrarily large constant M is added to the inter-cluster edge costs. CTSP does not constrain the order in which the clusters are visited. The clustered traveling salesperson problem with prespecified order (CTSP-PO) partitions the waypoints into clusters with specific priorities (Panchamgam, 2011). All waypoints in a given cluster must be visited before any waypoints in a lower-priority cluster are visited. Waypoint prioritization may be useful for MSAR because of limited UAV batteries. This prioritization will ensure that as many high-priority waypoints, assuming the available batteries are insufficient to fully cover the search area. Strict priority constraints may be beneficial if waypoint priority is critical; however, it may lead to extremely inefficient paths. Consider an MSAR search where there are two high probability regions: one at the top of a mountain and one at the bottom. Strict waypoint prioritization constraints will require that the UAV search these areas first before searching anywhere else on the mountain. This strategy is obviously inefficient when the UAV can easily search the area between the high-priority regions as it travels from one to the other. There exists a natural trade-off between path efficiency (i.e., length) and waypoint prioritization. The clustered traveling salesperson with d-relaxed priority rule (CTSP-d) parameterizes this trade-off with a parameter d (Panchamgam, 2011).

CTSP-PO and TSP are special cases of CTSP-d. CTSP-d becomes CTSP-PO when $d = 0$, and CTSP-d becomes TSP when $d = g - 1$, where g is the number of priority classes. The d-relaxed priority rule may provide the necessary control over the path efficiency/waypoint priority for MSAR UAV path planning. CTSP-d can be formulated as a mixed integer program (MIP) and solved optimally with tight worst-case bounds (Panchamgam et al., 2012). However, MIP-based solutions are not practical for in situ path planning. MIP Solvers (i.e., CPLEX, Gurobi) require substantial compute power and are not practical for applications like MSAR.

CTSP-d can be solved with a metaheursitic approach using a greedy randomized adaptive search procedure, iterative local search, and variable neighborhood descent with random neighborhood ordering (GILS-RVND) (Hà et al., 2022), and was first proposed to solve a minimum latency problem, which is a subset of the traveling salesperson problem (Silva et al., 2012). GILS-RVND can perform similarly to (i.e., 2% gap), and in some cases outperform the exact MIP solution, when the computation times were limited (Hà et al., 2022). The results demonstrate that GILS-RVND has a substantial computation speed advantage (approximately 10 minute solution time) with higher vertex counts (e.g., > 200). The authors provide recommendations for selection of certain parameters, which along with the specific mechanics of GILS-RVND are discussed in Section III.

The CTSP-d literature uses euclidean distance cost functions. Symmetrical cost functions are a natural choice for 2D domains, but are less applicable in 3D domains. Gravity creates cost (i.e., energy consumption) differences between UAV ascension and descension. An experimental study measured and modeled the energy costs of seven different UAV states: idle, armed, take off, hovering, horizontal flight, ascending, and descending (Abeywickrama et al., 2018). An improved cost function may incorporate the proportional energy consumption differences between different UAV states.

Most USA MSAR teams own UAVs, but lack the necessary software to plan and execute efficient, autonomous search coverage paths. The majority of existing coverage path planning algorithms are ill-suited for the particular MSAR UAV search requirements. Online algorithms are not practical due to the UAV's limited hardware computational capabilities. Many offline coverage path planning algorithms are too computationally expensive to be used on mobile hardware in remote environments. Decoupling viewpoint planning and path planning will minimize necessary in-field computing, while allowing MSAR personnel to plan UAV search paths with the most up-to-date information. The ability to prioritize high-probability search areas early in flights is essential for MSAR due to UAV's limited batteries;

however, the trade-off between path efficiency and search prioritization must be carefully selected. These factors must be addressed to bring the benefits of autonomy to MSAR.

III. PROPOSED SOLUTION

MSAR UAV search path planning requires 3D coverage path planning with scalable trade-off between viewpoint prioritization and path efficiency. Path planning must be tractable on portable hardware for in situ planning.

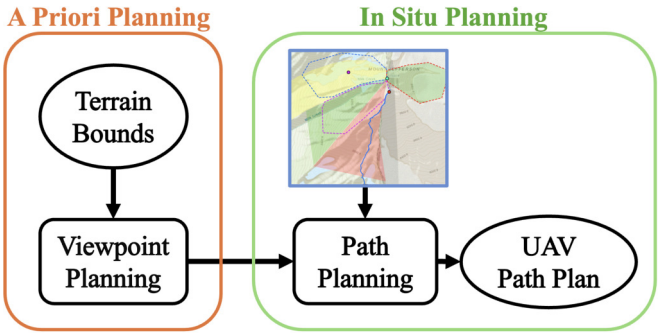

FIGURE 1: The proposed viewpoint planning and path planning steps. Viewpoints can be planned prior to a search, while UAV search paths can be planned in situ.

A two-step algorithm is proposed to address the identified limitations of existing methods that also meets MSAR requirements was developed as shown in Fig. 1. 3D sampling-based viewpoint planning with a novel sampling strategy efficiently generates viewpoints over a desired search area. The search area is defined using a probability map with discrete priority areas. A path between these viewpoints is planned using a metaheuristic CTSP-d solver with an energy-based cost function. The path planning algorithm controls the trade-off between viewpoint prioritization and path efficiency using a d-relaxation parameter. Decoupling viewpoint planning and path planning minimizes necessary in situ computation; thus, decreasing the compute costs experienced by MSAR personnel (e.g., battery requirements, computation time). It is infeasible to pre-plan UAV search paths for many MSAR scenarios due to the unreliability of information early in a search scenario. Paths must be planned in the field once MSAR personnel have established search areas and search area prioritizations.

A. VIEWPOINT PLANNING

Sampling-based viewpoint planning has been selected for its low computational requirements when compared to exact geometric methods. The selected method is similar to the voxel dilation method (Jing et al., 2019); however, this new approach uses independent sampling of latitude, longitude, and altitude to avoid expensive voxel dilation operations. Minimizing computational cost is important for MSAR, because the software must execute on common consumer hardware (e.g., laptops, tablets) within reasonable time.

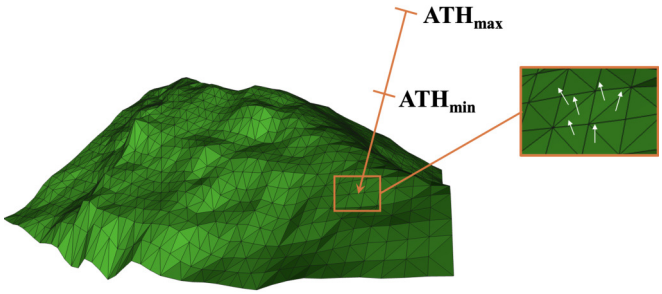

FIGURE 2: Viewpoints are generated by sampling a surface location within the search area, averaging surface normals to determine a view direction, projecting the average surface normal randomly between $\mathrm{ATH}_{\mathrm{min}}$ and $\mathrm{ATH}_{\mathrm{max}}$ meters above the surface, and adding noise.

Algorithm 1 Candidate Viewpoint Generation

Require: Terrain Bounds, $\mathrm{ATH}_{\mathrm{min}}$, Ground Sample Distance ($\mathrm{GSD}_{\mathrm{min}}$), Camera Field-of-View (FOV), Image Width in Pixels (I_x), Sample Density (S)

1: Download DEM data for the specified latitude/longitude bounds.
2: Generate a 3D mesh from DEM data.
3: N = surface normals for the triangles composing the mesh
4: $\mathrm{ATH}_{\mathrm{max}} = \frac{I_x}{2 \cdot \mathrm{GSD}_{\mathrm{min}} \cdot \tan(\frac{\mathrm{FOV}}{2})}$
5: $M = \mathrm{area}(\mathrm{Bounds}) * S$
6: $\mathcal{V} = \emptyset$.
7: $\mathcal{D} = \emptyset$.
8: **while** $|\mathcal{V}| < M$ **do**
9: $\quad x \sim \mathcal{U}(x_{\mathrm{min}}, x_{\mathrm{max}})$
10: $\quad y \sim \mathcal{U}(y_{\mathrm{min}}, y_{\mathrm{max}})$
11: $\quad z \sim \mathcal{U}(\mathrm{ATH}_{\mathrm{min}}, \mathrm{ATH}_{\mathrm{max}})$
12: $\quad vp'_i = \begin{pmatrix} x & y & z \times \mathbf{N}_{xy} \end{pmatrix}$
13: $\quad \mathrm{d}_i = -\mathbf{N}_{xy} + \sigma$
14: $\quad \mathcal{V} \leftarrow \mathcal{V} \cup \{v\}$.
15: $\quad \mathcal{D} \leftarrow \mathcal{D} \cup \{d\}$.
16: **end while**
return $\mathcal{V}$, $\mathcal{D}$

Viewpoint planning begins by downloading DEM data from the 3DEP dataset for a specified set of latitude/longitude bounds. A 3D mesh is generated from this raster using the Open3D library. Open3D calculates surface normals for the triangles composing the mesh. Candidate viewpoints are generated by sampling a latitude, longitude, and above ground level (AGL), as shown in Algorithm 1 lines 9, 10, and 11, respectively. $\mathcal{U}(a_{min}, b_{max})$ denotes the uniform distribution over the interval $[a, b]$. x_{min} and x_{max} define the latitude bounds, while y_{min} and y_{max} define the longitude bounds. The minimum above terrain height (ATH_{min}) is dictated by the desired minimum distance between the terrain and the UAV during flight. A lower minimum distance will decrease the area captured by each

image, increase the probability of terrain collisions, and increase the ground sample distance (GSD). The maximum ATH bound (ATH_{max}) is constrained by the desired GSD and the camera parameters (e.g., field-of-view, sensor size). ATH_{max} is calculated in Algorithm 1 line 4, where I_x is the image width in pixels, GSD_{min} is the minimum ground sample distance in pixels per meter, and FOV is field-of-view in degrees (Di Franco and Buttazzo, 2015).

Viewpoints at varying heights (between ATH_{min} and ATH_{max}) are planned to maximize coverage and minimize the number of viewpoints. A combination of viewpoints at varying heights will, on average, fully cover the terrain with fewer viewpoints than if all viewpoints were planned at a lower, constant above-ground height. Additionally, a higher, constant above-ground height will be more likely to not fully capture complex terrain, especially in the presence of obstacles (e.g., overhanging cliffs, small caves).

Candidate viewpoints (vp'_i) are calculated via Algorithm 1 line 12, where $\mathbf{N}_{xy}$ is the mean of all surface normals within the ATH_{min} distance of (x, y). This process, shown in Fig. 2, ensures the UAV maintains a minimum distance from the terrain. View direction (d_i) is calculated for each candidate viewpoint (i) by inverting the surface normal ($\mathbf{N}_{xy}$) and adding noise (σ), as shown in Algorithm 1 line 13.

The candidate viewpoints are down-sampled to a minimum coverage set. This set cover problem is solved using a greedy neighborhood search algorithm (Jing et al., 2019). These viewpoints can be saved for later path planning, assuming any desired search areas lie within the bounds of the viewpoints. Pre-planning viewpoints may reduce necessary in situ computation. Additionally, MSAR personnel may manually place additional viewpoints prior to a search to ensure coverage around known occlusions. MSAR personnel typically have substantial knowledge about the terrain in which they are operating. A path between viewpoints is planned once MSAR personnel establish a search area and prioritization map. The priority assigned to each viewpoint is equal to the highest priority region seen by the viewpoint.

MSAR personnel need a method of assigning priority to sviewpoint. This manuscript uses the SARTopo mapping software to define and export priority maps as GeoJSON files, shown in Fig. 3. GeoJSON is an open source geographic information system file format. SARTopo is frequently used throughout MSAR operations. Utilizing software already used by MSAR is an intuitive choice, as this will reduce the burden of learning additional technologies.

B. PATH PLANNING

The path between the selected viewpoints is planned via GILS-RVND with a desired d-relaxation value to control the trade-off between path efficiency and viewpoint prioritization (Hà et al., 2022). Two cost functions are used: the Euclidean distance (Euc.) (Equation 1) and a proxy energy model (Eng.) (Equation 2). The Euclidean distance cost function is the baseline approach, used by GILS-RVND (Hà et al., 2022). This cost function aims to minimize the total distance traveled. The proxy energy model cost function aims to penalize upward travel, which is more energy intensive than horizontal or downward travel. The proposed algorithm does not intend to limit path length to an estimated flight time of the vehicle at a certain velocity. The algorithm plans a coverage path for a given area, regardless of it's length. Prioritization of the high-probability regions of the search area maximizes search coverage if the path cannot be completed due to energy or other constraints.

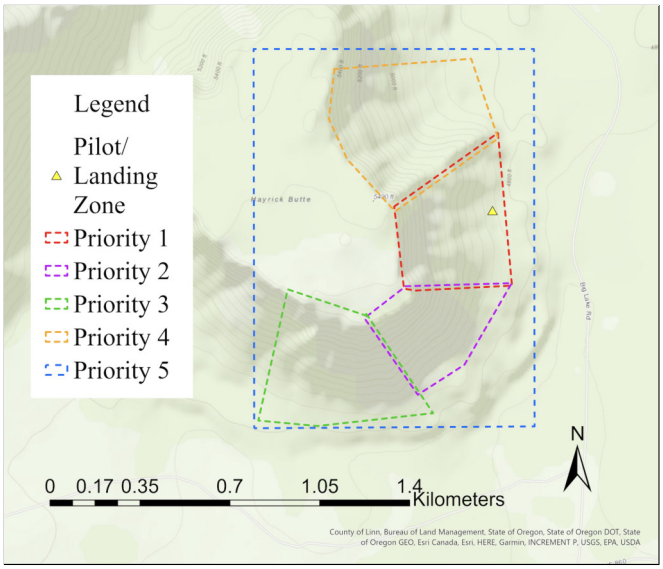

FIGURE 3: An example search area priority map is drawn in SARTopo for Hayrick Butte.

$$C_1 = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2 - z_1)^2} \quad (1)$$

$$C_2 = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2} + (z_2 - z_1) \quad (2)$$

IV. SIMULATION-BASED EXPERIMENTS

The combination of sample-based viewpoint planning and CTSP-d for UAV MSAR requires a series of experiments to demonstrate search efficacy and establish recommended parameters for different search scenarios. These algorithms have many parameters that must be carefully determined to yield the best planning results. Most MSAR personnel are not robotics experts. Viewpoint and path planning parameters must be constant values, selected procedurally, or very clear recommendations must be provided on how to select appropriate values. The primary objective of these experiments is to demonstrate the performance of the algorithms for UAV MSAR coverage search, and establish recommendations for how to configure the planners in different search scenarios. Table 1 provides the independent variables and values used for simulation-based experiments.

TABLE 1: The simulation-based experiment results were evaluated across a series of independent variables.

Variable	Values
Location	Hayrick Butte, North Sister, Mount Jefferson
ATH_{min} (m)	10, 40, 70
VP_{min}	1, 2
Cost Function	Euc., Eng.
D_{relax}	0, 2, 4
Path Plan Cutoff (minutes)	30

Minimum above terrain height (ATH_{min}) dictates the minimum distance, in meters, between the viewpoints and the terrain. The maximum above terrain height (ATH_{max}) is constrained by the desired ground sample distance and the camera configuration. $\text{ATH}_{max} = 83$ meters when line 4 of Algorithm 1 is used with the Parrot Anafi USA Camera Parameters and $\text{GSD}_{min} = 100$. $\text{ATH}_{max} = 83$ meters was rounded to $\text{ATH}_{max} = 80$ meters to provide a buffer against GPS positioning error. Lower values of VP_{min} will increase the probability of terrain collisions, but decrease the probability of visual terrain occlusions between the viewpoint and the terrain. A uniform distribution of values for ATH_{min} was chosen between 10 meters and 70 meters.

Minimum Viewpoints per patch (VP_{min}) determines how many viewpoints must capture each terrain mesh patch. An increase in VP_{min} will provide improved coverage of the terrain but increase UAV planning, energy, and data storage requirements. Euclidean and energy model cost functions are used for the path planning algorithm, as discussed in Section III. A path planning cutoff time of 30 minutes was chosen. Path planning times greater than 30 minutes will likely be burdensome to MSAR personnel; however, it is likely that some emergent scenarios will require less path planning time (i.e., 15 minutes). Each combination of these independent variable values was simulated for five iterations yielding a total of 540 simulations. The Parrot Anafi USA UAV configured with a dual-spectrum (i.e., visible, thermal) imager on a six degree-of-freedom gimbal is the assumed UAV for all field experiments.

The chosen dependent variables (i.e., metrics) are presented in Table 2. Viewpoint plan time and path plan time capture how long each of the planning steps take to execute. Number of viewpoints is the number of viewpoints selected by the viewpoint planner to completely cover a desired search area. Minimized numbers of viewpoints will minimize path planning computational complexity. Path Length and Elevation Gain measure length and elevation gain along the entire length of the planned path, capturing path efficiency. Ideally both of these variables are minimized for a given search area to minimize UAV energy consumption. Area coverage describes the percent of each priority area covered at a given point along the UAV path. These values will be 0% at the beginning of the path and 100% at the end.

TABLE 2: Dependent Variables.

Variable	Units
View Plan Time	minutes
Path Plan Time	minutes
Number of Viewpoints	count
Path Length	meters
Elevation Gain	meters
Area Coverage	percent

All simulations were performed on a 2023 Apple Macbook Air with M2 CPU and 16 gigabytes of RAM. This portable laptop is representative of computers that may be owned by MSAR personnel. A tablet computer is a more ideal candidate for in situ path planing; however, the Macbook Air has the same compute hardware as the 2023 Apple iPad Pro and is used as a proxy.

20 simulation trials were performed for each independent variable combination. Each simulation consisted of loading the DEM for the given location, constructing terrain mesh, planning viewpoints, and planning a UAV path. An MSAR subject matter expert drew five prioritized search areas for each of the evaluated locations: Hayrick Butte, North Sister, and Mount Jefferson.

Hayrick Butte is located in the central Oregon Cascade Mountains, and is frequented by backcountry skiers and snowmobile riders. MSAR personnel use the butte for training exercises due to its steep terrain, snow conditions, cliffs, and tree cover. Four search areas were priortized and drawn by a trained MSAR subject matter expert for a hypothetical scenario with a lost snowmobiler. A priority five search area (i.e., lowest priority) is the bounding box containing the other search areas. The resulting map is shown in Fig. 3, with the properties of the total search area provided in Table 3. The minimum Bound and Maximum Bound define the minimum and maximum latitude and longitude coordinates, respectively of the search area as universal transverse mercador (UTM) zone 10T coordinates. UTM is commonly used in MSAR. The depot, or landing zone, is the UAV's simulated take off and landing point that is selected to be an easily-accessible point in gentle terrain.

TABLE 3: Hayrick Butte Search Area Properties.

Minimum Bound	10T 0589889E 4917420N
Maximum Bound	10T 0590620E 4916374N
Depot	10T 0590491E 4916642N
Total Search Area	720 m^2
Min Elevation	1452 m
Max Elevation	1683 m

The North Sister scenario is based on a 2023 MSAR search when a climber fell while climbing to the summit

from the south. The MSAR subject matter expert, who was a member of a team that participated in this search, was asked to draw five prioritized search areas using information known at the beginning of the search (Fig. 4). This search area is larger and at a higher elevation in steeper terrain than Hayrick Butte, as shown by Table 4. Subjects who fall on North Sister often end up in dangerous locations that cannot be viewed or accessed safely by humans. UAVs will be particularly useful for identifying subject location and collecting information to inform a safe MSAR response.

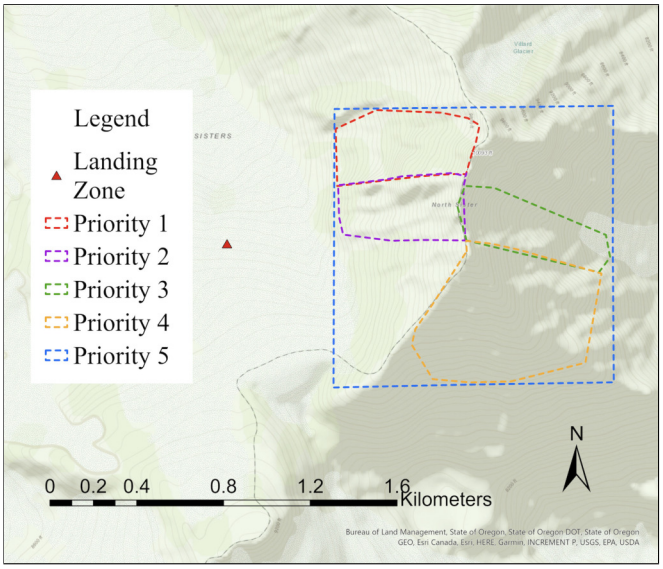

FIGURE 4: North Sister Search Area Map.

TABLE 4: North Sister Search Area Properties.

Minimum Bound	10T 0597700E 4891246N
Maximum Bound	10T 0598635E 4890358N
Depot	10T 0597357E 4890790N
Total Search Area	850 m^2
Min Elevation	2670 m
Max Elevation	3074 m

The Mount Jefferson scenario was based on the search as presented in Section 3. The MSAR subject matter expert participated in the actual search, but was instructed to draw a probability map with five priority areas using only knowledge from when the MSAR team was located at the Red Saddle on Mount Jefferson (Fig. 5), which is near the summit, and near where the subject was last seen. The Red Saddle may have been the most likely point for the MSAR team to deploy a UAV during the search. Mt. Jefferson is a useful test case for MSAR software, because it represents one of the most difficult locations in the Pacific Northwest to conduct an MSAR response. Mt. Jefferson is large (Table 5), extremely remote, has complex terrain with mixed conditions (e.g., dense trees, rocks, ice, snow), unpredictable weather, and can only be accessed on foot or via crewed aircraft.

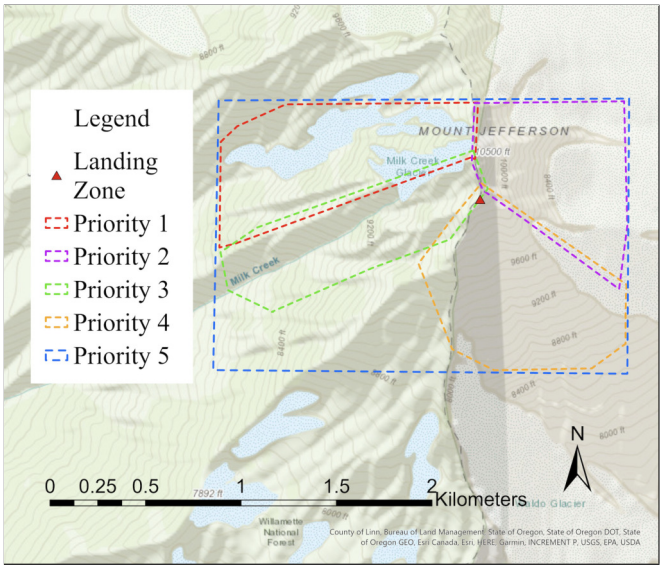

FIGURE 5: Mount Jefferson Search Area Map.

TABLE 5: Mount Jefferson Search Area Properties.

Latitude Bounds	10T 0594192E 4947668N
Longitude Bounds	10T 0595733E 4946677N
Depot	10T 0595171E 4947321N
Total Search Area	1560 m^2
Min Elevation	2377 m
Max Elevation	3201 m

A. RESULTS

Mean and standard deviation (in parentheses) will be provided for various metrics. The novel candidate viewpoint generation method was compared with the baseline voxel dilation method using viewpoint generation time. Both methods were used to generate one viewpoint for every 667 square meters, which is approximately four times the viewpoint density necessary to provide 98% coverage with 25% overlap, and a GSD of 33.04 pixels per meter. This viewpoint density was experimentally determined to yield sufficient coverage for the set coverage solver to select a subset. The new viewpoint generation method (Individual Sampling) is multiple orders of magnitude faster than the original voxel dilation method, as shown in Table 6.

Simulation results were evaluated for planning efficiency (i.e., time), path efficiency (i.e., length, elevation gain), and search efficiency. Search efficiency indicates how strictly the path follows viewpoint prioritization order. The effect of ATH_{min} and VP_{min} on path length were evaluated for all three locations. The mean path length results (shown in Table 7) demonstrate that path length decreases for Hayrick and

TABLE 6: Viewpoint sampling times (seconds) to achieve 98% coverage of search region, with minimum 25% adjacent image overlap for baseline and novel sampling methods.

Location	Voxel Dilation	Individual Sampling
Hayrick	48 (1.3)	0.06 (0.02)
North Sister	71 (1.9)	0.08 (0.01)
Jefferson	243 (0.32)	0.32 (0.04)

Mt. Jefferson, but increases for North Sister as the ATH_{min} increases. The ATH_{max} is a constant 85 meters, so increasing the ATH_{min} decreases the sample space; thus, reducing the overall distance between viewpoints. Additionally, increasing the VP_{min} from one to two increases the mean path length across all three locations, which is most obvious for Mt. Jefferson (Fig. 6). The Hayrick Butte path lengths ranged from 12.1 kilometers (km) to 19.2 km, the North Sister path lengths ranged from 17.2 km to 27.5 km, and the Mt. Jefferson path lengths ranged from 29.5 km to 50.0 km.

TABLE 7: Path length (kilometers - km) descriptive statistics for various viewpoint planning parameters.

ATH_{min}	10		40		70	
VP_{min}	1	2	1	2	1	2
Hayrick	14.2 (0.7)	18.2 (0.8)	13.4 (0.7)	17.2 (0.8)	13.0 (0.7)	16.2 (0.6)
North Sister	20.9 (1.0)	26.1 (0.7)	18.6 (0.7)	24.3 (0.6)	19.1 (0.8)	24.9 (0.7)
Mt. Jefferson	32.9 (1.0)	43.7 (1.3)	31.2 (0.7)	41.6 (1.1)	31.2 (1.0)	40.6 (1.1)

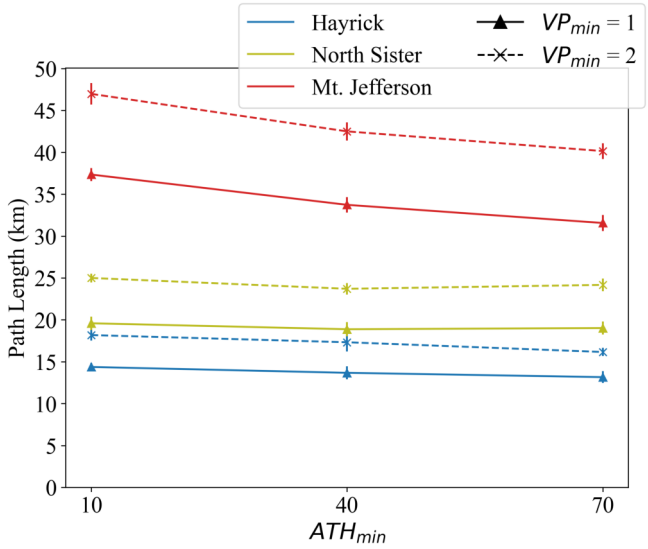

FIGURE 6: Path length by ATH_{min}, VP_{min}, and location.

Mean elevation gain decreases as the ATH_{min} increases, with relatively small standard deviations (Table 8), which is reinforced by Fig. 7. This relationship is a product of the reduced viewpoint sample space as the ATH_{min} increases. The maximum vertical distance between two viewpoints is equal to ATH_{max} minus ATH_{min}. This maximum distance is minimized when $\text{ATH}_{min} = 70$. Increasing the VP_{min} increases the mean elevation gain, with the largest effect for the Mt. Jefferson results. This result is likely caused by the search area size. Doubling the required number viewpoints substantially increases the total vertical distance between viewpoints, especially when the search area is large.

TABLE 8: Total elevation gain (m) descriptive statistics for various viewpoint planning parameters.

ATH_{min}	10		40		70	
VP_{min}	1	2	1	2	1	2
Hayrick	2.2 (0.1)	3.1 (0.2)	2.0 (0.1)	2.7 (0.1)	1.7 (0.1)	2.2 (0.1)
North Sister	5.0 (0.3)	6.7 (0.3)	4.3 (0.2)	5.6 (0.2)	4.2 (0.2)	5.4 (0.2)
Mt. Jefferson	7.3 (0.3)	10.5 (0.3)	6.8 (0.2)	9.5 (0.3)	6.6 (0.3)	8.7 (0.3)

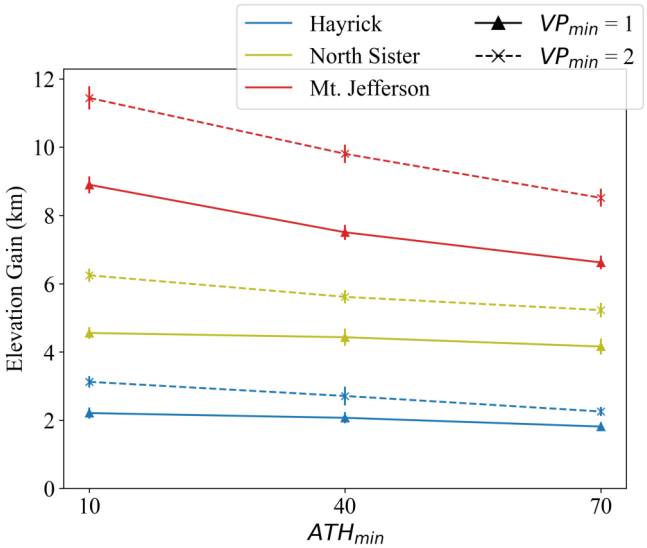

FIGURE 7: The mean elevation gain by ATH_{min} and VP_{min}.

Viewpoint planning time is highly reliant on the number of viewpoints being sampled and selected, which is a direct result of search area size. The mean viewpoint planning time is primarily a product of the search location, as seen in Table 9. The ATH_{min} and VP_{min} have no obvious impact on viewpoint planning time, as shown by Fig. 8. The viewpoint planner is not affected by the vertical distance between candidate viewpoints (i.e., when ATH_{min} increases), or the

total number of viewpoints being selected from a given candidate set (i.e., when VP_{min} increases).

TABLE 9: View planning time (minutes) descriptive statistics for various view planning parameters.

ATH_{min}	10		40		70	
VP_{min}	1	2	1	2	1	2
Hayrick	1.81 (0.04)	1.91 (0.01)	1.94 (0.02)	1.98 (0.02)	1.93 (0.01)	1.97 (0.01)
North Sister	2.46 (0.07)	2.44 (0.02)	2.43 (0.07)	2.53 (0.05)	2.42 (0.02)	2.49 (0.06)
Mt. Jefferson	13.12 (2.52)	11.91 (0.17)	11.83 (0.08)	12.25 (0.17)	11.70 (0.03)	12.49 (0.09)

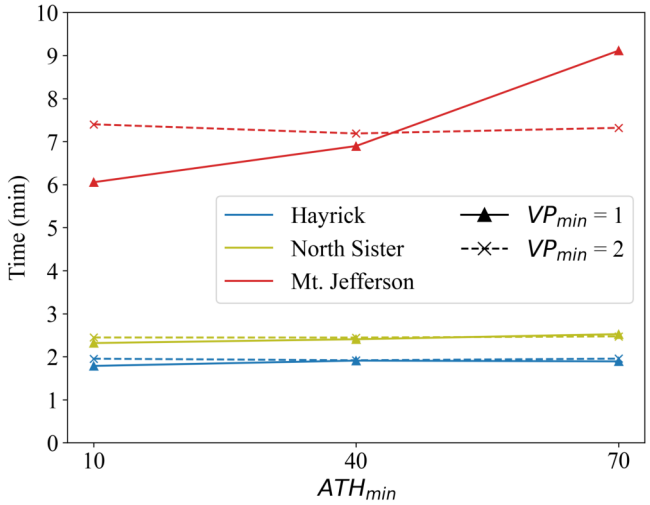

FIGURE 8: Viewpoint plan time by ATH_{min} and VP_{min}.

The Hayrick and Mt. Jefferson simulations demonstrated a decrease in number of planned viewpoints as ATH_{min} increased (Table 10). The North Sister results experienced variation in the number of viewpoints with increased ATH_{min}, which is likely a product of terrain complexity. North Sister has the most complex terrain of the three locations. Increasing the ATH_{min} may impair occlusion avoidance; therefore, more viewpoints are required to capture the terrain. This explanation is reinforced by the extra increase of viewpoint count for North Sister where $ATH_{min} = 70$ and $VP_{min} = 2$, as shown by Fig. 9. Increasing the number of viewpoints required to observe North Sister's complex terrain further increased the viewpoints required.

Mean path length was calculated for all combinations of d-relaxation, cost function, and location (Table 11). The mean path length was relatively unchanged between the two cost functions, which is visible in Fig. 10. The mean path length decreased as the d-relaxation parameter increased, for both Hayrick and North Sister. The mean Mt. Jefferson simulation path lengths increased slightly when the d-relaxation parameter increased from zero to two, but decreased when

TABLE 10: Descriptive statistics for the number of planned viewpoints by ATH_{min} and VP_{min}.

ATH_{min}	10		40		70	
VP_{min}	1	2	1	2	1	2
Hayrick	172.1 (2.73)	295.5 (5.48)	157.2 (3.87)	265.9 (7.32)	150.2 (2.43)	243.6 (1.00)
North Sister	256.5 (11.71)	414.4 (8.95)	220.4 (2.68)	377.0 (5.85)	231.7 (2.47)	411.9 (10.81)
Mt. Jefferson	400.1 (1.36)	683.3 (2.71)	374.0 (0.0)	629.1 (4.97)	375.0 (0.0)	644.6 (0.90)

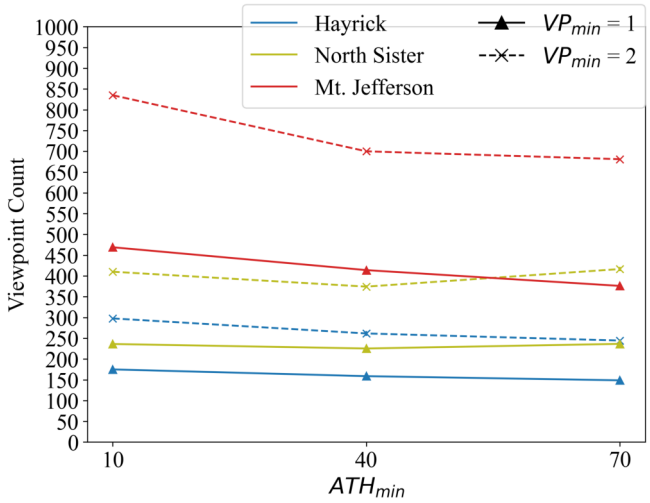

FIGURE 9: Mean viewpoint count by ATH_{min} and VP_{min}.

d-relaxation was increased, but these results are within one standard deviation, as shown by Fig. 10.

TABLE 11: Path length (km) descriptive statistics for d-relaxation and cost function parameters.

d-Relaxation	0		2		4	
Cost Function	Euc.	Eng.	Euc.	Eng.	Euc.	Eng.
Hayrick	16.2 (2.0)	16.2 (1.9)	15.4 (2.1)	15.4 (2.1)	14.8 (1.9)	14.8 (2.0)
North Sister	22.9 (2.9)	22.9 (2.8)	22.3 (2.9)	22.3 (3.0)	21.8 (3.1)	21.7 (3.1)
Mt. Jefferson	37.0 (5.1)	36.9 (4.9)	37.6 (5.4)	37.5 (5.6)	36.1 (5.3)	36.1 (5.4)

The mean elevation gain results reflect the same relationship as seen for path length, given the independent variables. The mean elevation gain for Hayrick and North Sister generally decrease; albeit, within one standard deviation (Table 12). The mean elevation gain for Mt. Jefferson had a much higher standard deviation with no clear relationship to the d-relaxation value (Fig. 11). The cost function had no discernible effect on mean elevation gain across all locations.

Path planning was limited to 30 minutes. There is a direct relationship between the search locations and planning time

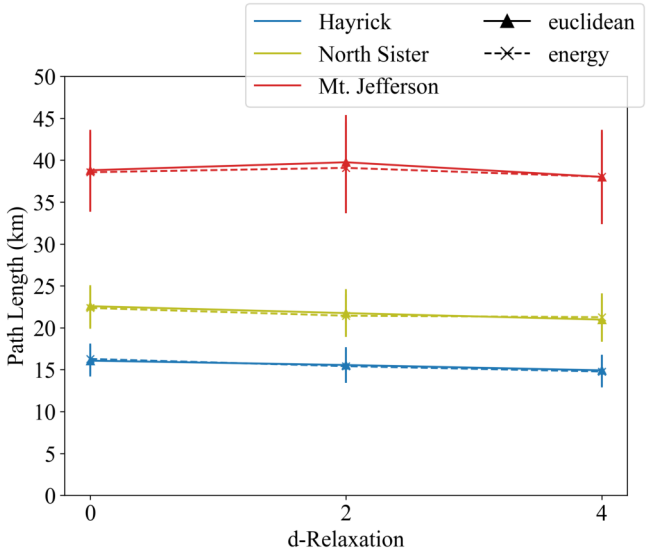

FIGURE 10: The mean path length by d-relaxation parameter, cost function, and location.

TABLE 12: Elevation gain (kilometers) descriptive statistics for d-relaxation parameter and cost functions.

d-Relaxation	**0**		**2**		**4**	
Cost Function	**Euc.**	**Eng.**	**Euc.**	**Eng.**	**Euc.**	**Eng.**
Hayrick	2.5 (0.5)	2.5 (0.5)	2.3 (0.5)	2.4 (0.5)	2.3 (0.5)	2.3 (0.5)
North Sister	5.3 (0.9)	5.3 (0.9)	5.2 (0.9)	5.2 (0.9)	5.1 (0.9)	5.0 (0.9)
Mt. Jefferson	8.2 (1.4)	8.2 (1.4)	8.4 (1.5)	8.5 (1.5)	8.1 (1.5)	8.0 (1.5)

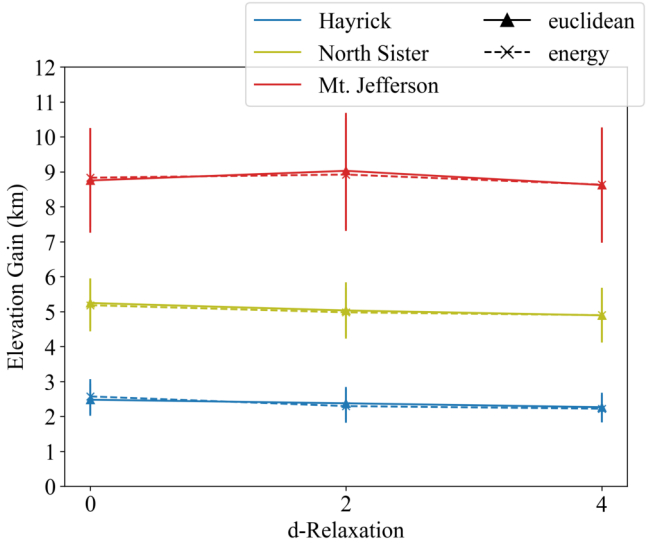

FIGURE 11: The mean elevation gain by d-relaxation parameter, and the two cost functions.

(Table 13). Mt. Jefferson is the largest search area and has the greatest mean path planning time across all parameters. The mean path planning times of greater than the 30 minute threshold are a result of necessary overhead processes that must occur before the planner executes (e.g., loading viewpoints into memory, assigning viewpoint priorities). The large standard deviations for path planning times of Hayrick Butte simulations and the North Sister simulation when d-relaxation parameter is zero, indicate that planning time is highly variable when it completes before the 30 minute threshold (Fig. 12). All the Mt. Jefferson simulations and the North Sister simulations for the two and four d-relaxation parameters were terminated at the 30 minute cutoff. The CTSP-d solver is an anytime algorithm so any path planning run that is terminated early yields a viable path; however, it will likely be sub-optimal.

TABLE 13: Descriptive statistics for path planning time (minutes) by d-relaxation parameter and cost function.

d-Relaxation	**0**		**2**		**4**	
Cost Function	**Euc.**	**Eng.**	**Euc.**	**Eng.**	**Euc.**	**Eng.**
Hayrick	17.1 (9.4)	17.5 (9.0)	22.6 (7.5)	23.0 (7.3)	27.4 (4.5)	27.3 (4.6)
North Sister	25.5 (5.4)	25.4 (5.5)	30.2 (0.1)	30.1 (0.3)	30.3 (0.1)	30.2 (0.1)
Mt. Jefferson	30.2 (0.3)	30.2 (0.4)	30.6 (0.5)	30.5 (0.3)	30.9 (0.8)	31.0 (0.9)

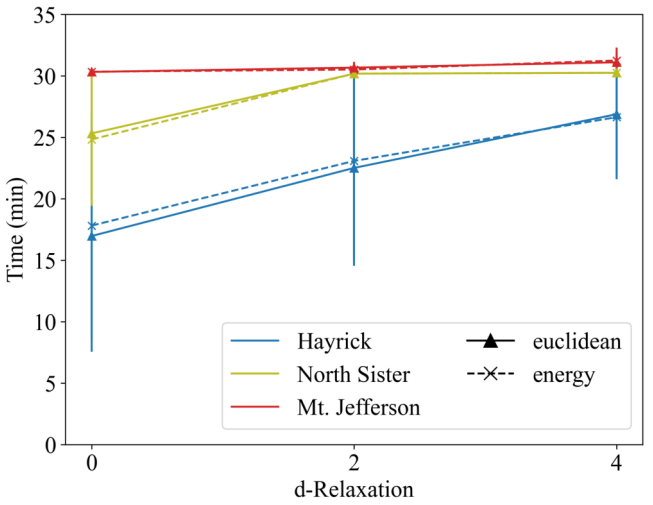

FIGURE 12: The mean path planning time by d-relaxation parameter and the two cost functions.

B. DISCUSSION

The relationship between terrain complexity and ATH_{min}'s effect on viewpoint count must be considered when selecting an appropriate value for ATH_{min}. Simple terrain (e.g., Hayrick Butte) appear to benefit from a higher value of

ATH_{min}. The simulation results indicate that more complex terrain (e.g., North Sister, Mt. Jefferson) can have highly variable outcomes based on the ATH_{min} values. A programmatic method of selecting this value may be ideal; however, it may be difficult to capture terrain complexity quantitatively in a way that can be used to pragmatically select an appropriate ATH_{min} value.

The cost function results indicate that the euclidean cost function sufficiently captures the relative cost of traveling between viewpoints, which is likely a product of the pre-planned viewpoints. The viewpoints are planned without considering the cost (i.e., length, elevation gain) of the resultant path. The simple euclidean cost function is likely sufficient for path planning because the vertical travel cost between viewpoints is relatively low, compared to the horizontal travel cost between viewpoints. Path planning with a higher viewpoint density, where the the vertical and horizontal distance between adjacent viewpoints is similar, may benefit from the energy cost function.

Total altitude gain may be decreased by dynamically increasing the viewpoint sample space. The algorithm samples a statically defined viewpoint planning space defined by the ATH_{min} and ATH_{max} values. Increasing ATH_{min}, thus decreasing the sample space, will decrease path efficiency. An improved viewpoint sampling strategy may sample a set of viewpoints with ATH_{min} equal to ATH_{max}. The set cover problem can be solved for these viewpoints. An iterative sample-then-solve process can decrease ATH_{min} with each step. This approach may minimize the total elevation gain by maximizing the coverage at each sample step.

Path length and elevation gain are expected to decrease as the d-relaxation value approaches $N-1$, where N is the number of prioritized search areas. The planner increasingly prioritizes path efficiency (i.e., length, elevation gain) over viewpoint prioritization, as the d-relaxation value increases. This relationship is clear in the Hayrick Butte results, of lesser magnitude in the North Sister results, and non-existent in the Mt. Jefferson results. These behaviors are observed for both cost functions. This effect is likely due to the increased number of viewpoints (a product of the search area size) relative to the number of prioritized search areas. Consider an edge case where viewpoints are randomly distributed, and viewpoints within each priority class have no spatial order. Changes to the d-relaxation parameter will have substantial impact on the efficiency of the resultant path. A path planned with a d-relaxation parameter of zero may be substantially longer than a path planned with a d-relaxation parameter of $N-1$. The antithetical edge case occurs when the viewpoints within each priority class are spatially clustered. A d-relaxation parameter value of zero represents the worst-case for path efficiency; however, in this case, it will likely not impact overall path efficiency, because efficient paths can be planned within each cluster.

The GILS-RVND solver is an anytime algorithm that allows early termination; however, early termination of the iterative CTSP-d will, on average, result in a less efficient path. The Hayrick Butte results indicate that higher d-relaxation parameters require longer computation times. It follows that a path plan with a d-relaxation parameter of four terminated at 30 minutes will be less efficient than a path plan with d-relaxation parameter of zero terminated at the same cutoff time. The Hayrick Butte results demonstrate that when the path planning processes are run to completion (e.g., less than 30 minutes), a higher d-relaxation value yields a shorter path. This behavior is expected, as higher d-relaxation values prioritize path efficiency over viewpoint prioritization. There was no mean path length reduction for Mt. Jefferson with a increased d-relaxation value, and no mean path length reduction for North Sister when the d-relaxation parameter was greater than two and all trials reached the 30 minute planning cutoff time.

The simulation results can be used to develop suggestions for viewpoint and path planning parameters. A large value of ATH_{min} (e.g., ATH_{min} = $\text{ATH}_{min} - 15$ meters) will provide the most efficient paths for simple terrain. Complex terrain may benefit from a lower ATH_{min}, but specific suggestions are difficult to derive from the varied simulation results. Increasing VP_{min} greatly decreases path efficiency. This parameter must be minimized (e.g., $\text{VP}_{min} = 1$), unless there is a specific, situational need for increased viewpoints (e.g., many occlusions not captured by the DEM). The euclidean cost function will likely provide more efficient paths than the energy cost function. Lastly, the d-relaxation parameter did not have a substantial impact on the planning time; as such, minimizing this parameter will likely benefit MSAR path planning. Following strict priority will maximize search efficiency and provide the highest probability of subject detection if the path is terminated early.

V. FIELD EXPERIMENT

Field experiments were performed to evaluate the system's real-world performance, and answer the question- how does viewpoint prioritization effect the real-world execution of search paths planned in-situ for MSAR? The field experiment conditions were complete snow coverage, sparse evergreen trees, and cliffs near the top of the butte, as shown in Fig. 13. The snow cover and tree density is shown in Fig. 14. The field experiments were conducted at Hayrick Butte over the course of three days in February 2024. The Parrot Anafi USA UAV is small enough to carry in a backpack, such as the ones shown in Fig. 14. This UAV has a stated flight time of 32 minutes under ideal environmental conditions (e.g., calm winds, moderate temperatures, sea level). This UAV model was selected because it supports programmatic control of the UAV via the Parrot Olympe software development kit, which is a Python library that provides a high-level UAV control interface. This software was used to control the UAV and collect imagery. Multiple members of the Corvallis Mountain Rescue Unit participated in the field testing, including the

MSAR subject matter expert who drew the Hayrick Butte search area priority map, shown in Fig. 3.

The viewpoints were planned prior to the field experiments and paths were planned in the field. The field experiment independent variables are provided in Table 14. $\text{ATH}_{min} = 70$ was chosen because of Hayrick Butte's simple terrain, and ATH_{max} was chosen to maximize the coverage from each viewpoint. The euclidean cost function was chosen, because the simulation results indicated that the energy cost function provided no benefit. A d-relaxation of zero was chosen to prioritize search efficiency over path efficiency, since the simulation results indicated that smaller d-relaxation values do not have a large negative impact on path efficiency. Search priority areas for the field exercise were the same as those used in the simulation experiments, as shown in Fig. 3.

FIGURE 13: The UAV taking off from the Hayrick Butte field experiment depot during field experimentation.

FIGURE 14: All gear must be carried by humans.

TABLE 14: The field experiment independent variables.

Variable	Values
Location	Hayrick Butte
(ATH_{min}, ATH_{max})	(70, 85), (55, 70)
Minimum Viewpoints Per Patch	1
Cost Function	Euc.
d-Relaxation	0
Path Plan Cutoff (minutes)	30

The paths were evaluated by analyzing the percent area covered of each priority area against the percent of path completed. This analysis provides insight into the prioritization of certain areas relative to other areas. This analysis is particularly useful for paths that are terminated early, demonstrating what percentage of each path was complete.

Seven flights were conducted over three days. The first day was dedicated to site selection and initial flight testing. Selection of an appropriate landing zone is critical for the safety and efficacy of UAV deployments in backcountry environments. The landing zone must be free of significant collision risks (e.g., trees, rocks, artificial structures), have a level surface for takeoff/landing, and be within line-of-sight of the UAV pilot. Additionally, the landing zone must be located near the intended search area. Spatial separation from the search area will increase required battery capacity and decrease the amount of time the UAV can search.

The landing zone was placed within search area one and as near the base of Hayrick as possible. The landing pad dug into the snow was approximately 0.5m^2 in size. Snow conditions, such as those shown in Fig. 15, greatly impact UAV take-off and landing operations. Time and care must be taken to sufficiently prepare a landing zone. The selected landing zone was located within the first clearing of sufficient size with good visibility of the intended search areas. Placing the landing zone low on the butte limited required hiking time to the landing zone, which is especially important for time critical MSAR operations.

FIGURE 15: A UAV launch area dug into the snow.

The first day's test flights were designed to prove the functionality of safety features, like the return-to-land (RTL) behavior when the path was complete, or the UAV's battery reached the low battery threshold. The second and third days were dedicated to debugging and experimentation.

The weather data was acquired from the Santiam Pass Ski Patrol seven day weather summary website (Patrol, 2024). The air temperature was favorable throughout the field experiments, with the temperature varying from -7^oC to -2^oC, as shown by Fig. 16. This temperature range was well above the Parrot Anafi USA UAV's stated -35.5^oC minimum operating temperature (Parrot, 2022).

Wind conditions were favorable for the first two days, as shown in Fig. 17 and 18. The wind conditions on 02 February

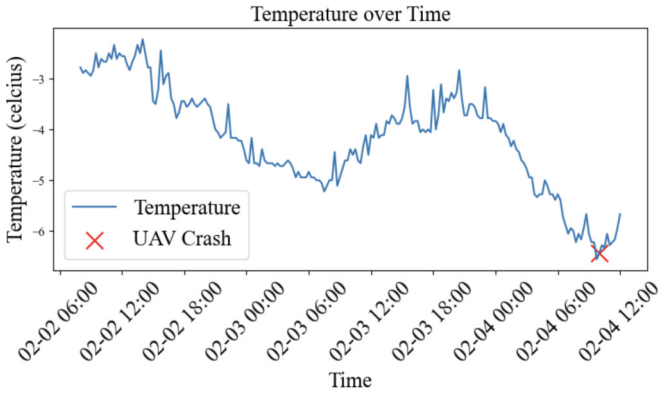

FIGURE 16: Air temperature data during the field trials acquired from the Santiam Pass Ski Patrol weather station located at the Hoodoo Ski Area.

at 17:30 changed significantly when the wind changed from a gentle south westerly wind to a strong easterly wind. The final flight performed on 04 February at 09:18 resulted in lost line-of-sight, followed by a presumed UAV crash. The autopilot was terminated and RTL was commanded, but the UAV did not return. A search of Hayrick Butte near the final viewpoint did not successfully recover the UAV. There were no environmental indications of high winds at the landing zone on 04 February, because the landing zone was sheltered via trees and terrain. Once the UAV crested the top of Hayrick on the final flight, it was caught in the strong easterly wind and was unable to RTL.

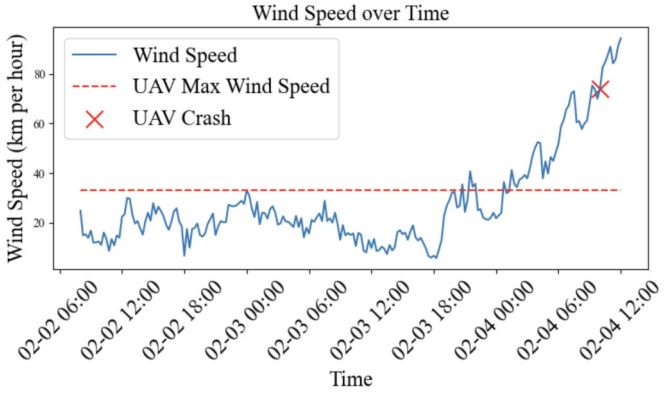

FIGURE 17: The Santiam Pass Ski Patrol weather station's wind data during the field trials.

A. RESULTS

Seven UAV flights were performed with varying success. The UAV was commanded to takeoff to 30 meters above the landing zone, and fly to each viewpoint, with the UAV's and gimbal's poses dictated by the planned view direction. The UAV paused at each viewpoint to capture one image before proceeding to the next viewpoint. The pilot terminated the autopilot and manually commanded RTL if UAV line-of-

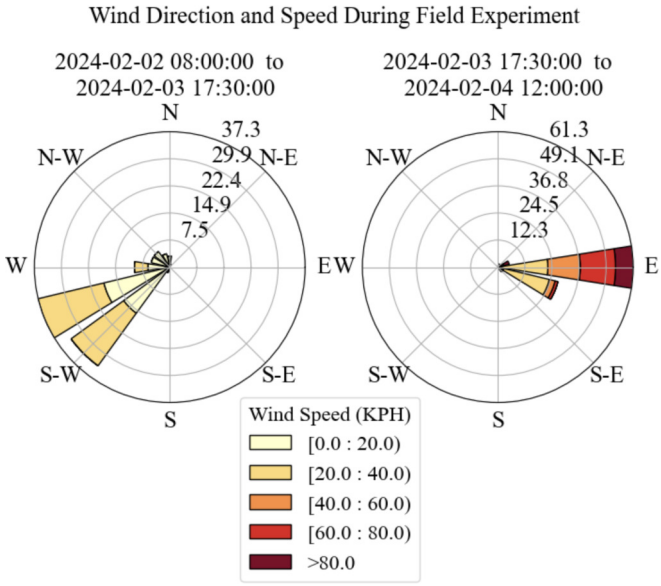

FIGURE 18: The wind direction and magnitude changes.

sight was lost. Five flights were terminated due to lost line-of-sight, one was terminated due to a cold battery, and one was terminated due to reaching the low battery threshold. Maintaining line-of-sight was the responsibility of the pilot. Additional visual observers were tasked with maintaining line-of-sight with the UAV for some of the flights. Initially, all visual observers were located at the landing zone.

The first path was planned to perform strict priority coverage (i.e., d-relaxation = 0) using the same prioritized search map as the Hayrick Butte simulations (Fig. 19). Path 1's parameters are provided in Table 15. VP_{min}, Cost Function, and d-Relaxation were all constant for these paths with values of 1, euclidean, and 0, respectively. The number of viewpoints, path length, and elevation gain of the resultant path fall within the expected values, based on the simulation results (Table 16). The path planning completed in 7.5 minutes, which is well below the 30 minute threshold.

TABLE 15: Parameters for the three field testing paths.

Name	Map	ATH_{min} (m)	ATH_{max} (m)
Path 1	Original	70	85
Path 2	Original	55	70
Path 3	Reduced	55	70

TABLE 16: Path metrics for the three paths.

Name	N Viewpoints	Path Length (km)	Elevation Gain (km)	Path Plan Time (min)
Path 1	154	13.59	1.95	7.50
Path 2	171	13.56	1.96	10.42
Path 3	55	4.25	0.79	0.57

The first two UAV flights used Path 1, shown in Fig. 19 with a maximum flight speed of 5 meters per second (m/s). This low flight speed was initially used to increase safety. The first two flights (i.e., 03 Feb 10:24, 03 Feb 10:39) were canceled due to lost line-of-sight between the pilot and the UAV. All visual observers were located at the landing zone. Line-of-sight was lost because the ATH_{max} of 85 meters was too high to maintain line-of-sight. The UAV's small size relative to the distance from the pilot, combined with the bright sky and glare from the snow made maintaining line-of-sight difficult. The first flight completed 32% of the 13.59 km path, as shown by Fig. 20. This flight covered 94% of Priority 1, 92% of the Priority 2, and 94% of Priority 3 areas.

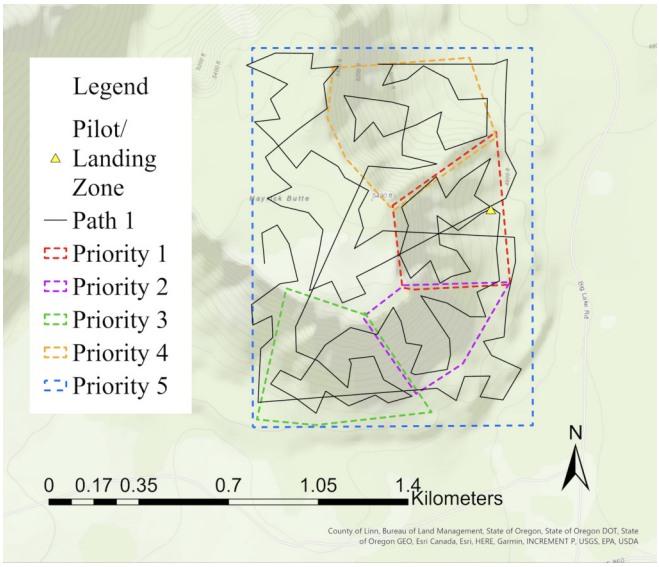

FIGURE 19: Path 1 planned with original Hayrick map.

Line-of-sight during the 03 Feb 10:39 flight was lost due to terrain occlusion. Approximately 90% of the priority 1 search area was covered, and 80% of the Priority 2 search area were covered in less than 20% of the total flight path length, before return-to-land was commanded (Fig. 21).

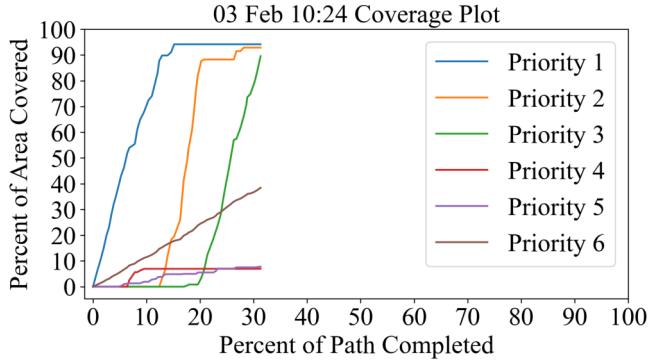

FIGURE 20: 03 Feb 10:24 priority area coverage.

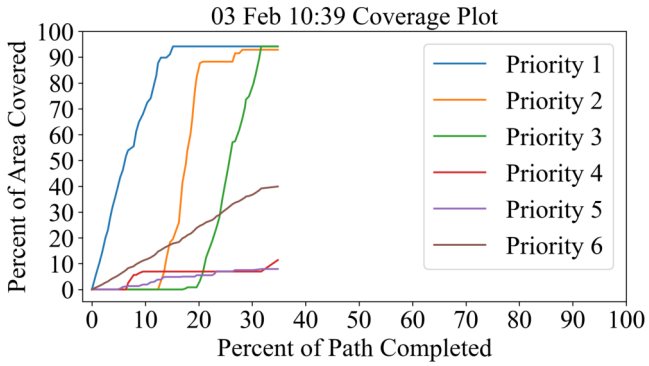

FIGURE 21: 03 Feb 10:39 priority area coverage.

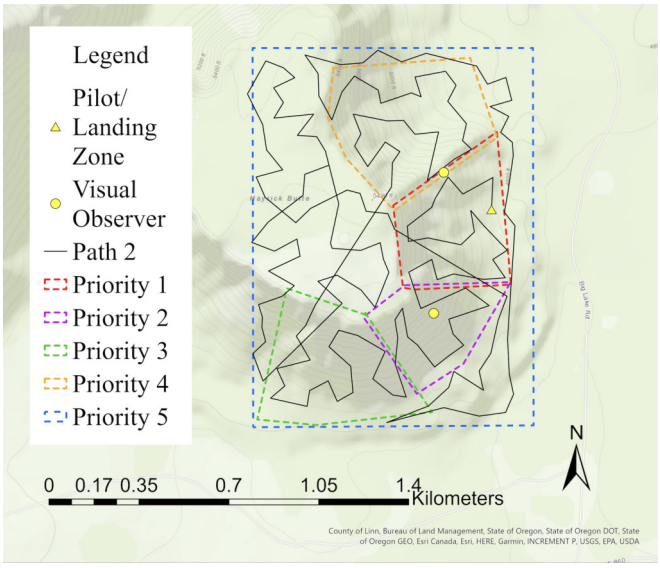

FIGURE 22: Path 2 was planned using the same priority map as path 1, but with a reduced viewpoint sample space.

Path 2 was planned with the ATH_{min} and ATH_{max} values reduced by 15 meters in an attempt to improve line-of-sight to the UAV, as shown by Table 15. Decreasing the altitude of the viewpoint sample space increased the number of viewpoints, but did not have a large impact on path length, or elevation gain (Table 16). The path planning completed in approximately 10 minutes, which was longer than Path 1, but below the 30 minute threshold.

Two visual observers moved to search areas two and three to improve UAV line-of-sight, as shown by Fig. 19. Each observer indicated via radio when they acquired or lost UAV line-of-sight. The next two flights (i.e., 03 Feb 11:07, 03 Feb 12:14) were flown with a maximum flight speed of 10 m/s.

The 03 Feb 11:07 flight was canceled due to lost line-of-sight, but achieved 100% coverage in priority areas 1 and 2, 99% in priority areas 3 and 4, and 82% in priority area 5, yielding a total coverage ratio of 92% (Fig. 23).

The 03 Feb 12:14 flight was terminated early by a low temperature battery error. The UAV will automatically land

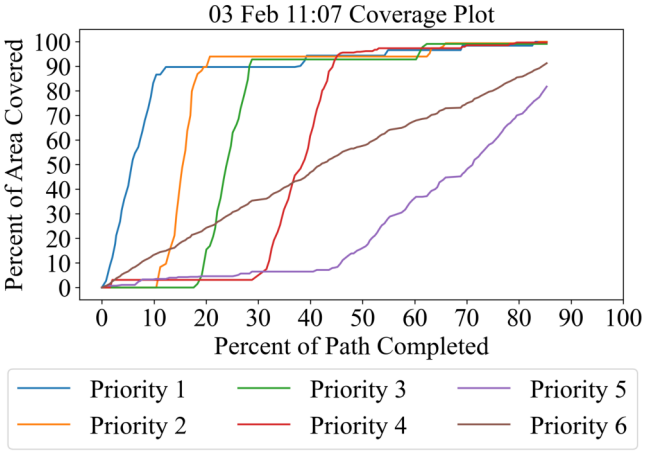

FIGURE 23: 03 Feb 11:07 priority area coverage.

if the battery temperature is too low. The battery used in this flight was previously sitting in a backpack, hence the low temperature. This flight covered 90% of the Priority 1 search area, and 82% of the Priority 2 search area before the emergency landing (Fig. 24). The UAV auto-landed into a tree high on Hayrick Butte, but was recovered.

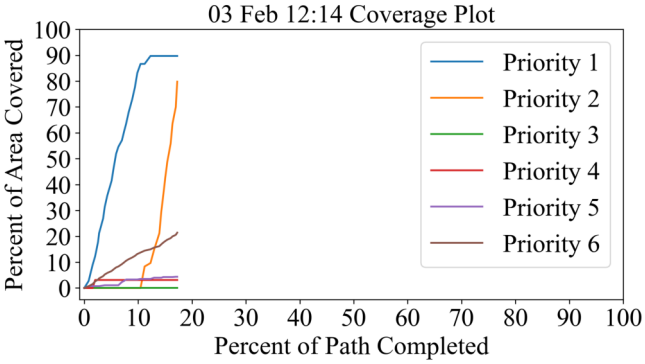

FIGURE 24: 03 Feb 12:14 priority area coverage.

The next two flights (i.e., 03 Feb 13:23, 03 Feb 13:50) were flown with a maximum flight speed of 15 m/s. The 03 Feb 13:23 flight covered 90% of the priority 1 search area, 94% of the priority 2 search area, and 93% of the priority 3 search area, yielding a total area coverage of 49%, as shown in Fig. 25. This flight was terminated due to lost line-of-sight with the UAV. The 03 Feb 13:50 flight covered 97% of the priority 1 and 4 search areas, 99% of the priority 2 and 3 search areas, and 45% of the priority 5 search area, yielding a total coverage of 73%, as shown in Fig. 26. This flight was terminated when line-of-sight with the UAV was lost.

The MSAR subject matter expert was asked to draw a new prioritized search area map that was constrained to terrain visible by the visual observers, as shown in Fig. 27. The properties for this new search are are shown in Table 17. A third path was planned using the reduced prioritized search

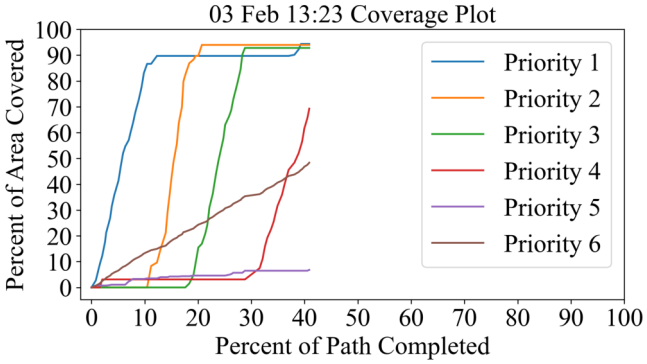

FIGURE 25: 03 Feb 13:23 priority area coverage.

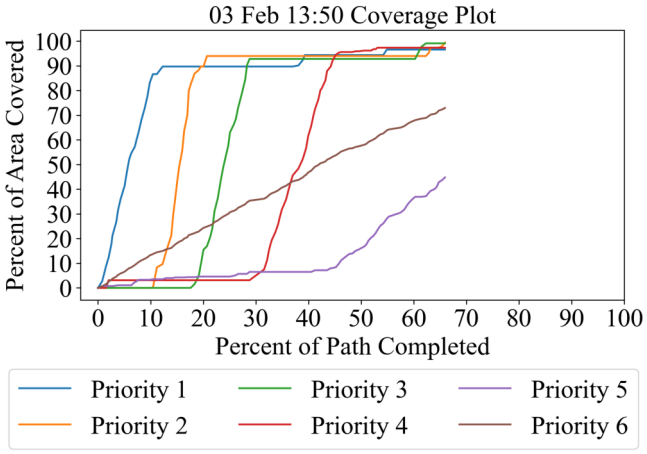

FIGURE 26: 03 Feb 13:50 priority area coverage.

area map. Path 3 was planned with the same parameters as Path 2, but for this new map. This reduction in search area decreased the number of viewpoints by 58%, the path length by 69%, the elevation gain by 59%, and path planning time by 95%, as compared to Path 2 (16). The flight covered 90% of search area Priority 1 and 70% of search area Priority 2 before line-of-sight was lost. Three visual observers were located at the landing zone, one was located in the priority 2 search area, and one was located in the priority 4 search area. The pilot commanded a return-to-land, but, the UAV did not return, likely due to the drastic wind changes that were not detectable at the launch area.

TABLE 17: Hayrick Butte Reduced Search Area Properties.

Minimum Bound	10T 0590155E 4916334N
Maximum Bound	10T 0590557E 4916918N
Depot	10T 0590491E 4916642N
Total Search Area	230 m^2
Min Elevation	1462 m
Max Elevation	1672 m

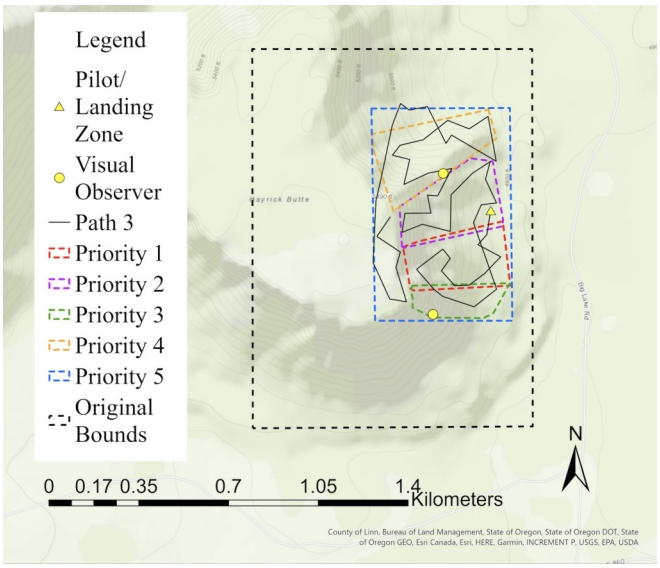

FIGURE 27: Search area reduced to improve line-of-sight.

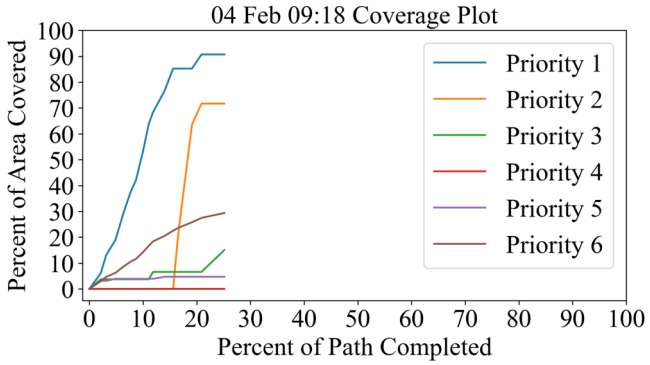

FIGURE 28: 04 Feb 09:18 priority area coverage.

B. DISCUSSION

Each flight resulted in covering at least 90% of Priority area 1 and often covered greater than 90% of the other priority areas, even though none of the seven flights completed the planned paths. This coverage was achieved at 10-15% path completion for all seven flights. Five of the flights covered at least 90% of the three highest priority search areas at or before 30% path completion. These results indicate that the prioritized search strategy can improve coverage search in unfavorable conditions (e.g., poor visibility, high winds). The strict prioritization provided high coverage of the highest priority search areas, when compared to the total path length. Coverage of the highest priority search areas happened early in the flight paths, which is critical for MSAR UAV operations to ensure that if a flight is terminated early, the path will have covered as much of the high-priority areas as possible. This strategy is critical for MSAR where path completion is unlikely due to many compounding factors (e.g., weather, battery problems, visibility constraints). Additionally, searching the highest-priority areas first minimizes subject detection time, if data is streamed back to the ground searchers and processed in real time. The proposed algorithm planned paths in the field well below the 30 minute threshold; however the terrain at Hayrick butte is less complex than other locations, as demonstrated in the simulation based results for North Sister and Mt. Jefferson.

Operating UAVs in the field is known to be difficult due to a variety of compounding factors (e.g., battery problems, line-of-sight requirements, weather conditions). The problems faced throughout these field experiments are a reality of MSAR UAV operations. The initial field experiment objective was to fly complete paths and analyze the priority area coverage across those flights. Instead, the flights were plagued by the same issues that have been, and will continue to be experienced by MSAR teams. The coverage of the highest priority areas early in the short flights indicate that this algorithm can allow MSAR teams to collect useful search data, even if the full path is not completed.

VI. DISCUSSION

The simulation results for viewpoint candidate generation demonstrated that the novel viewpoint sampling strategy is significantly more computationally efficient than the baseline method. This new approach does not sacrifice coverage, because regardless of the viewpoint sampling strategy used, the viewpoints are down-sampled to a minimal coverage set by the set coverage problem solver. This decrease in computational requirements is especially vital, because these algorithms must execute on any consumer device within reasonable time bounds. This will additionally allow generation of viewpoints for larger search regions.

The simulation results demonstrated that increasing the ATH_{min} can improve path efficiency (i.e., decreased path length, decreased total elevation gain) for some terrain, but have the opposite effect for other terrain. This result is likely a product of terrain complexity. The North Sister simulation results showed that a moderate value for the ATH_{min} provided optimal coverage with the least number of viewpoints. The field experimentation results demonstrated that the ATH_{max} is constrained by the camera parameters, the desired ground sample distance, and the visibility limitations. The ATH_{max} must be set below the cloud ceiling. The field experiments demonstrated that if the $\text{ATH}_{max} > 70$ in sunny, snowy conditions, it was difficult to maintain line-of-sight of the UAV. This value will vary depending on environmental conditions and the deployed UAV. The specific constraints of MSAR (e.g., personnel carrying all equipment, limited resupply) necessitate small form-factor UAVs; therefore, UAVs used by MSAR are likely to be similar form factor to the UAV used in these field experiments.

The complex effect of ATH_{min} on path efficiency makes it difficult provide clear guidelines for choosing a value for all possible MSAR situations. The best path efficiency in simple terrain (e.g., Hayrick Butte) will result from a small

viewpoint sample space (e.g., $\text{ATH}_{max} - \text{ATH}_{min} = 15$ meters). There may exist a procedural method to select an optimal ATH_{min} value for more complex terrains. The viewpoint planning algorithm's current objective is to minimize the total number of viewpoints; however, the total number of viewpoints selected is inconsequential. Minimizing path length and total elevation gain is the true objective. These metrics are difficult to calculate without solving a path at each step of the viewpoint planning process, which will likely be intractable on a consumer-grade computer.

The proposed energy-based cost function did not provide any obvious benefit during the simulation-based experiments; however, this cost function may be used to design an improved objective function for the viewpoint planning algorithm. Currently, the viewpoint planning algorithm iteratively selects the next best viewpoint based on which adjacent viewpoint increases coverage the most. Changing the next viewpoint selection process to find the viewpoint that maximizes the coverage and minimizes the elevation change relative to the horizontal travel distance may yield more efficient paths for all terrain types.

$\text{VP}_{min} = 1$ will likely be optimal for most MSAR scenarios. $\text{VP}_{min} = 1$ provides full area coverage with a set percentage overlap between adjacent images. Simulation results indicate that $\text{VP}_{min} = 2$ greatly decreases path efficiency, thus decreasing coverage for a given UAV energy availability. This trade-off may be beneficial if the search area contains many occlusions not captured by the DEM (e.g., large avalanche debris field, boulders from a landslide). Future research can consider how to select separate VP_{min} values for each priority area, which can provide improved coverage in the high priority search areas without greatly impacting overall path efficiency. Additionally, enforcing a minimum view angle difference between two viewpoints of one terrain patch may further improve coverage around occlusions. Two viewpoints will likely not provide improved coverage of a mesh patch if their view directions are parallel.

The simulation results indicate that d-relaxation does not have have a substantial impact on path efficiency; thus, lower d-relaxation values (i.e., valuing viewpoint prioritization over path efficiency) will likely benefit UAV MSAR operations. This fact is reinforced by the field experiment. Every flight was terminated early due to some problem (e.g., cold battery, low battery, lost line-of-sight, weather); however, because all paths were planned with a d-relaxation value of zero, every flight covered at least 90% of the priority 1 search area.

Image-based MSAR search is a challenging task that suffers from the unfortunate reality that humans must search the collected images. The proposed algorithm improves the imagery collection process, but will still require humans to manually search the imagery. Ideally, computer vision can be leveraged to perform real-time object detection of search clues; however, computational limitations of consumer-grade UAVs will likely prevent on-board processing. The computer or tablet used for path planning may also be used for real-time object detection on imagery streamed from the UAV.

It is important to note that many of the planned paths are infeasible to fly with a single battery (i.e., less than 32 minutes at 5m/s); however, the path planning algorithm does not constrain the overall path length. The algorithm minimizes path length given search prioritization constraints. The resultant paths may require one or more batteries; however, this fact highlights the benefit of this approach. There is no reliable method of calculating path feasibility and the total flight time available for MSAR operators deploying a UAV for search will be highly dependent on many factors (e.g., available batteries, environmental conditions). Therefore, planned paths must prioritize high-probability regions early in the flight to maximum search efficiency for a given constant, however unknown, flight time.

Future research must consider using other sensing modalities (e.g., avalanche transceivers, Recco) on the UAV to permit real-time actionable information. These sensors provide continuous data that does not follow the same paradigm as individual image collection (i.e., a single image is collected at each viewpoint); therefore, the inter-viewpoint paths are valuable sensing opportunities. Path primitive-based planning plans with continuous path segments, rather than individual viewpoints, may better leverage sensors providing continuous data (Jing et al., 2019). Paths planned with discrete viewpoints for continuous data sensors does will not plan for the inter-viewpoint data collection; thus, unnecessarily increasing the sensor coverage.

Additionally, the algorithm can be improved by adding an additional planning step to guarantee sensor coverage. The simulation and field experiments presented in this manuscript achieved 100% terrain coverage; however this is not explicitly guaranteed by the algorithm, and may fail for more complex terrain. A second viewpoint planning step may procedurally insert viewpoints to cover any uncovered regions. Lastly, the algorithm may be improved by adding inter-viewpoint collision checks. The experiments presented in this manuscript did not suffer from any collisions; however this is a vital feature for real-world MSAR applications.

The simulation and field experiment results demonstrate that priority-aware search planning can improve search efficiency when path completion is uncertain. The computationally efficient viewpoint and path planning algorithms will allow MSAR personnel to plan coverage paths on consumer-grade hardware without a large time burden. A reasonably efficient path can be planned in situ, in 30 minutes or less, for search areas of similar size to those used in the simulation experiments. Increasing the search area size for a given path planning cutoff time will increase the efficiency gap between the solution and the optimal solution. Path efficiency will also be reduced if the path planning time limit or compute power is reduced. This approach can also be beneficial to other areas of robotics. Autonomous inspection and coverage planning, especially in constrained or austere

environments (e.g., nuclear reactors, spacecraft inspection), are active research areas. Task completion uncertainty is a reality of field robotics, and algorithms that prioritize high-priority areas will provide resilience against this uncertainty.

VII. CONCLUSION

MSAR is a challenging domain for robotics, due to many compounding constraints (e.g., computation limitations, environment). Adoption of UAVs among MSAR teams demonstrates that teams are willing to overcome the challenges of UAV adoption (e.g., initial cost, training, regulations) to gain safety and potentially improve search outcomes. MSAR teams have yet to fully realize the benefit of UAVs, due to available path planning software not meeting the specific requirements of MSAR (e.g., offline, computationally efficient). This manuscript presents a novel approach to this problem by introducing a computationally efficient method of 3D viewpoint planning with the use of CTSP-d to allow viewpoint prioritization in the planning process. The proposed solution blurs the line between hasty or comprehensive search. Currently, UAVs for MSAR are used with two separate strategies: hasty or comprehensive search. These disjoint strategies are inefficient for MSAR where UAV flight time is limited by the batteries carried by rescuers. Rescuers cannot afford to fly separate, hasty and comprehensive search flights. Simulation experiments indicated that because the d-relaxation parameter had no significant impact on path efficiency, a d-relaxation value of 0 (i.e., strict-prioritization) is likely optimal for MSAR search, as this provides the highest search coverage if a path is incomplete.

The reality of MSAR operations is that search paths will often be incomplete due to energy limitations, environmental conditions, and line-of-sight constraints. Future research must investigate improved heuristics for selecting d-relaxation values, especially when the search priority map is less uniformly clustered (i.e., two distinct areas of equal priority within an area of lesser priority). Future research may explore the feasibility of planning based on visual observer locations to better maintain line-of-sight.

REFERENCES

Abeywickrama, H. V., Jayawickrama, B. A., He, Y., and Dutkiewicz, E. (2018). Comprehensive Energy Consumption Model for Unmanned Aerial Vehicles, Based on Empirical Studies of Battery Performance. *IEEE Access*, 6:58383–58394. Conference Name: IEEE Access.

Bolourian, N. and Hammad, A. (2020). LiDAR-equipped UAV path planning considering potential locations of defects for bridge inspection. *Automation in Construction*, 117:103250.

Chisman, J. A. (1975). The clustered traveling salesman problem. *Computers and Operations Research*, 2(2):115–119.

Choset, H. and Pignon, P. (1998). Coverage Path Planning: The Boustrophedon Cellular Decomposition. In Zelinsky, A., editor, *Field and Service Robotics*, pages 203–209. Springer London, London.

Daudelin, J. and Campbell, M. (2017). An adaptable, probabilistic, next-best view algorithm for reconstruction of unknown 3-d objects. *IEEE Robotics and Automation Letters*, 2(3):1540–1547.

Deboodt, T. (2023). Corvallis Mountain Rescue Unit. Corvallis, Oregon. 3 Feb 2024 (personal communication).

Di Franco, C. and Buttazzo, G. (2015). Energy-aware coverage path planning of uavs. In *2015 IEEE International Conference on Autonomous Robot Systems and Competitions*, pages 111–117.

Goodrich, M. and Lin, L. (2009). UAV Intelligent Path Planning for Wilderness Search and Rescue. *Faculty Publications*.

Goodrich, M. A., Morse, B. S., Gerhardt, D., Cooper, J. L., Quigley, M., Adams, J. A., and Humphrey, C. (2008). Supporting wilderness search and rescue using a camera-equipped mini UAV. *Journal of Field Robotics*, 25(1-2):89–110.

Hà, M. H., Nguyen Phuong, H., Tran Ngoc Nhat, H., Langevin, A., and Trépanier, M. (2022). Solving the clustered traveling salesman problem with -relaxed priority rule. *International Transactions in Operational Research*, 29(2):837–853.

Jing, W., Deng, D., Xiao, Z., Liu, Y., and Shimada, K. (2019). Coverage path planning using path primitive sampling and primitive coverage graph for visual inspection. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1472–1479.

Jing, W., Polden, J., Lin, W., and Shimada, K. (2016). Sampling-based view planning for 3d visual coverage task with unmanned aerial vehicle. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1808–1815.

Jordan, S., Moore, J., Hovet, S., Box, J., Perry, J., Kirsche, K., Lewis, D., and Tse, Z. T. H. (2018). State-of-the-art technologies for UAV inspections. *IET Radar, Sonar & Navigation*, 12(2):151–164.

Maboudi, M., Homaei, M., Song, S., Malihi, S., Saadatseresht, M., and Gerke, M. (2023). A review on viewpoints and path planning for UAV-based 3-d reconstruction. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16:5026–5048.

Manerba, D., Mansini, R., and Riera-Ledesma, J. (2017). The traveling purchaser problem and its variants. *European Journal of Operational Research*, 259(1):1–18.

Matai, R., Singh, S., and Mittal, M. (2010a). *Traveling Salesman Problem: an Overview of Applications, Formulations, and Solution Approaches*.

Matai, R., Singh, S. P., and Mittal, M. L. (2010b). Traveling salesman problem: an overview of applications, formulations, and solution approaches. *Traveling salesman problem, theory and applications*, 1(1):1–25.

Muchiri, G. N. and Kimathi, S. (2022). A Review of Applications and Potential Applications of UAV. *Proceedings of the Sustainable Research and Innovation Conference*, pages 280–283.

Nam, T. H., Shim, J. H., and Cho, Y. I. (2017). A 2.5d map-based mobile robot localization via cooperation of aerial and ground robots. *Sensors*, 17(12).

Panchamgam, K., Xiong, Y., Golden, B., Dussault, B., and Wasil, E. (2012). The hierarchical traveling salesman problem. *Optimization Letters*, 7(7):1517–1524.

Panchamgam, K. V. (2011). *Essays in retail operations and humanitarian logistics*. University of Maryland, College Park.

Parrot (2022). Parrot anafi usa white paper. https://www.parrot.com/assets/s3fs-public/2022-01/white-paper-anafi-usa-v1.5.3_en.pdf.

Patrol, S. P. S. (2024). Hoodoo butte weather seven day summary. http://www.santiampassskipatrol.org/wx/weather3.htm.

Polidori, L. and El Hage, M. (2020). Digital elevation model quality assessment methods: A critical review. *Remote Sensing*, 12(21).

Santini, A., Viana, A., Klimentova, X., and Pedroso, J. P. (2022). The probabilistic travelling salesman problem with crowdsourcing. *Computers and Operations Research*, 142:105722.

Schomer, N. and Adams, J. A. (2024). Mountain search and recovery: An unmanned aerial vehicle deployment case study and analysis. [In Review].

Schomer, N. L. and Adams, J. A. (2022). Unmanned Aerial Vehicle Usage and Affordability Among Mountain Search and Rescue Teams. In *IEEE International Symposium on Safety, Security, and Rescue Robotics*, pages 61–66.

Silva, M. M., Subramanian, A., Vidal, T., and Ochi, L. S. (2012). A simple and effective metaheuristic for the minimum latency problem. *European Journal of Operational Research*, 221(3):513–520.

Skydio (2023). Skydio 3D Scan. https://www.skydio.com/3d-scan.

Stoker, J. and Miller, B. (2022). The accuracy and consistency of 3d elevation program data: A systematic analysis. *Remote Sensing*, 14(4).

Tan, C. S., Mohd-Mokhtar, R., and Arshad, M. R. (2021). A Comprehensive Review of Coverage Path Planning in Robotics Using Classical and Heuristic Algorithms. *IEEE Access*, 9:119310–119342.

Vansteenwegen, P., Souffriau, W., and Oudheusden, D. V. (2011). The orienteering problem: A survey. *European Journal of Operational Research*, 209(1):1–10.

Vincent-Lambert, C., Pretorius, A., and Van Tonder, B. (2023). Use of Unmanned Aerial Vehicles in Wilderness Search and Rescue Operations: A Scoping Review. *Wilderness & Environmental Medicine*, 34(4):580–588.

Whitehead, K. and Hugenholtz, C. H. (2014). Remote sensing of the environment with small unmanned aircraft systems (UASs), part 1: A review of progress and challenges. *Journal of Unmanned Vehicle Systems*, 02(03):69–85.