Recursive Non-Uniform Coverage of Unknown Terrains for UAVs

Seyed Abbas Sadat, Jens Wawerla and Richard T. Vaughan
Autonomy Lab, Simon Fraser University
{sas21, jwawerla, vaughan}@sfu.ca

***Abstract*— Area coverage is one of the compelling applications of UAVs. The existing methods for coverage path planning assume a uniformly interesting target area. However in many real world applications items of interest are not uniformly distributed but form clusters. Here it can be advantageous to only sample regions of interest while skipping uninteresting sections of the environment. In this paper, we present a coverage tree structure that can accommodate non-uniform coverage of regions in the target area. Three strategies are proposed to traverse the coverage tree. Experiments indicate that in some situations our method can cover the interesting regions with about half the travel time / cost of a naive regular 'lawnmower' coverage pattern.**

I. INTRODUCTION

UAVs equipped with light-weight sensors are becoming affordable platforms for fast data acquisition with many applications. One particular task in which UAVs are used is area coverage in applications like agriculture [1], surveillance [2], vegetation monitoring [3], terrain mapping [4], etc. Informally speaking, in an area coverage task the robot must move along a path so that the footprint of a specific sensor on the robot sweeps the whole area of interest.

Many algorithms have been developed to solve the coverage path planning problem [5]. In all these methods, the area to be covered is assumed to be uniformly interesting and consequently the robot moves within a specific distance from the surface. However, in many applications, as a result of non-uniformity in the environment (Fig. 1), different parts of the target area can be covered with different resolutions, for example, by flying at different altitudes. This may allow the path planner to produce shorter paths due to the fact that the sensor footprint sweeps a bigger area as the distance between the sensor and the target surface increases [6]. In many real-world applications the distribution of interesting sub-areas is not known in advance. But we may be able to use the available data to classify a section as possibly interesting or uninteresting. At high altitude we can decide online whether there is a need to cover a sub-region from closer viewpoints or not. Based on this capability, we use a coverage tree structure that can accommodate non-uniform coverage of different regions in the target area. Three different strategies are introduced to cover the area by traversing the coverage tree. All strategies are complete, i.e. they will cover the target area up to the required resolution for each sub-region. At the same time, each strategy has features which might make it the better choice in different situations.

Simulation experiments were used to compare the efficiency of these strategies with the standard 'lawnmower' coverage pattern of uniform stripes. We also performed two simulations based on real-world vegetation data; a target application for the method.

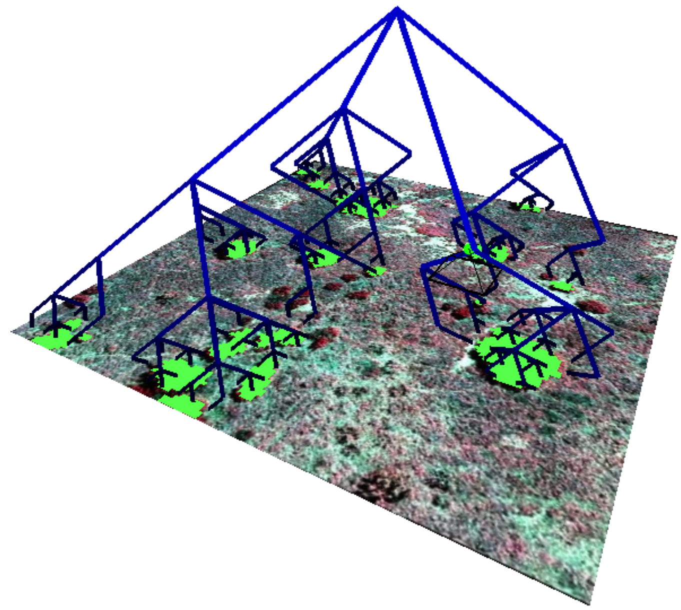

Fig. 1: Simulated environment for area coverage with UAV. Interesting regions are colored green. The lines show the interesting branches of the coverage tree.

The rest of this paper is organized as follows: related research is presented in the next section followed by section III in which the structure of the coverage tree is described. In section IV the strategies to traverse this tree are proposed. Experiments and results are discussed in section V followed by the conclusion and future work in section VI.

II. RELATED WORK

UAVs have been used for aerial imaging in many projects. A single quad-rotor is used in [7] to cover an irregular area. The user specifies the requirements (such as resolution, image overlapping, etc.) of the images as well as the target area and the system plans and executes a coverage path. In [1] multiple UAVs are used to take georeferenced images of farm land. In order to create a full map of the area using image mosaicking techniques, grid-based coverage path planning was used to fully cover the target area [8]. Polygon decomposition is used in [9] to allocate members of a UAV team to different sub-regions based on individual capabilities. For each sub-region, the sweeping direction of the lawnmower pattern is selected to minimize the number of turns. Unmanned helicopters were used in [10] for automatic crop dusting. Simple back-and-forth motions are used to cover segments of the field after decomposition. A team of hex-rotors is used in [11] to take high quality images from farm fields by visiting a predefined set of waypoints. The images are then used by agricultural experts to locate weed pods.

Seabed coverage using an autonomous under-water vehicle (AUV), which resembles aerial coverage with UAVs, has also been studied in recent years. In [12], an AUV is used to map and visualize a large region of seafloor using the high-resolution images captured by the camera on the robot. The high cost of movement in the vertical axis is considered in [13] to generate coverage paths for non-planar regions of seafloor. Galceran et. al. in [6], noticed that in lawnmower-like seafloor coverage, an AUV that stays at a constant depth, will lead to undesirable coverage overlap among the back-and-forth laps caused by the change in seafloor height. To alleviate this problem, they segmented the area into regions of constant height and generated lawnmower sweeping patterns with different inter-lap spacing and sweeping direction. Their simulations showed that the generated coverage paths were shorter than those of previous methods. The idea of covering the regions of interest is also present in [14]. However, the interesting regions are extracted offline from previous lawnmower-like surveys, in contrast to our approach which is online with no prior survey.

A recent review of coverage path planning in general can be found in [5]. Many approaches generate coverage paths by decomposing the target area (with possible obstacles) into simple convex regions and generating simple lawnmower-like paths to sweep each region. In [15], multi-robot coverage in terrains with non-uniform traversability is studied. Uniform coverage of structures with complex topology is studied in [16] motivated by automotive spray painting.

The contributions of this paper are methods for online coverage path planning that consider non-uniformity in the unknown environment. The proposed methods can replace the uniform lawnmower-like patterns (that are almost always used to cover simple regions) to achieve shorter coverage paths. Consequently, interesting areas will be covered with higher resolution sensor coverage than less interesting areas.

III. COVERAGE TREES

Let us assume that the target area A is $m \times m$ meters and free of obstacles. Also, for simplicity, assume that the shape of the sensor footprint is a square with length of $l(h)$ where h is the distance of the sensor to the ground. The coverage tree embedded in the metric space is recursively created as follows:

i The root of the tree R, is located at the center of A with a height of $h_R = l^{-1}(m)$ (where $l^{-1}(.)$ is the inverse of the function $l(.)$), i.e. the sensor footprint covers A at the root node.

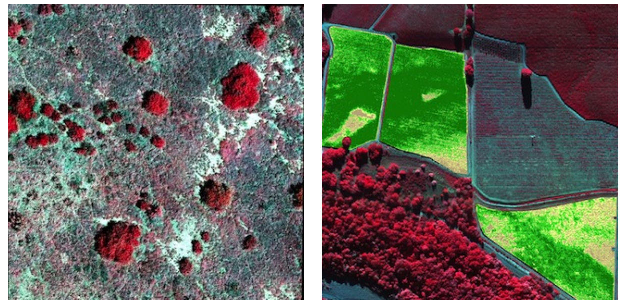

(a) Searching for trees in a desert-like environment

(b) Covering fields of particular crop.

Fig. 2: Real environments used in our simulations

ii Let h_n and A_n be the height of node n and the area covered by the sensor at node n respectively. Then, for a *branching factor*[1] $b \geq 2$, A_n is decomposed into a $b \times b$ grid with cells of length $l_c = \frac{l(h_n)}{b}$ and therefore $h_c = l^{-1}(l_c)$. For some threshold h_t, if $h_c \geq h_t$ then there is a node for each grid cell, with node n as parent, at the center of the cell and with height h_c, i.e. the sensor covers the grid cell at the child node.

According to the above definition, for a $m \times m$ area A, the root node will be at the center and at a height such that the whole area is covered by the footprint. Then, assuming $b = 2$, the root node will have 4 child nodes forming a 2×2 grid, at lower height $h^{'} = l^{-1}(\frac{m}{2})$ so that each child node covers $\frac{1}{4}$ of A and so on.

In coverage trees, the child nodes cover exactly the same area as the parent does but with a higher resolution which depends on the branching factor b. Also no nodes with a height less than a threshold h_t exist. The parameter h_t determines the lowest height at which the sensor can sufficiently cover a region at the finest resolution. Therefore regular path planners for area coverage use this height to produce a lawnmower pattern (assuming no obstacles in the area) to sweep the whole area. This simple pattern is equivalent to visiting all the leaf nodes of the coverage tree (in an order that yields the minimum path length).

If we know that a region covered by the children of a node is not interesting, it is enough (and possibly more efficient) to cover that region with lower-resolution sensing by only visiting the parent node. Since this information in not available *a priori*, one can visit the parent node first and then decide whether or not the sub-regions covered by each child node are interesting enough to be visited. We assume that the sensor data can always correctly indicate which parts of the footprint are interesting and need closer coverage, i.e., our interestingness sensor has no false negatives. For example, images capture by a camera can be processed to recognize some particular vegetation that makes the green parts in Figure (2b) interesting and the red parts uninteresting.

Here we try to keep the notion of interestingness as general as possible. For instance, there does not have to be an explicit

[1] *Branching factor* is usually used to refer to the number of children of each parent in a tree. However, here we use it to denote the parameter b.

part of the sensor data that shows the interestingness of a region directly. It can be the result of complex probabilistic reasoning that leads to the need for closer coverage of a region. For example, if the goal of the coverage is to find people in an unknown area, regions with buildings/structures are more interesting than bushes, and as the UAV descends in the coverage tree, direct people recognition can be used as the interestingness metric. Similarly, when the UAV is mapping the underlying terrain, regions with more uncertainty or local variation in height can have priority for closer coverage than the rest of the area. In case of false positives on interestingness, path length may be increased but coverage is still guaranteed.

The branching factor b controls the rate of the increase in resolution of the coverage between a parent node and its children. A large branching factor will produce a coverage tree composed of nodes with coarse coverage at high altitudes and children that are close to the surface of the area. This configuration of the tree will be practical if the estimation of the interestingness is very accurate. Otherwise a smaller branching factor will be more viable.

The shape of the sensor footprint is assumed to be square to simplify the systematic decomposition of the area. If the footprint of a sensor is in another shape e.g. circle, then one can use the largest square that fits inside the footprint and use it as the square footprint assumed in coverage trees. For instance, when the sensor is a camera with wide angle of view lens, the largest square in the middle of the undistorted image is used as the effective footprint (see for example [17]). Both the sensor footprint and the target area can have an arbitrary shape and the only requirement is a method to decompose the area into a number of footprint shapes with minimum overlap. Everything else about the coverage tree remains the same.

In case of a rectangular target area A, we construct the coverage tree for the smallest square that contains A. Then we prune every node k for which $A_k \cap A = \emptyset$, i.e. the sensor footprint at node k does not intersect with the target area. Therefore we can accommodate non-square areas as well.

Note that if the environment has obstacles which prevent the UAV from flying close to the surface, the corresponding nodes will not be feasible and can be removed from the tree. In this way we can easily relax the assumption that no obstacles are present in the area.

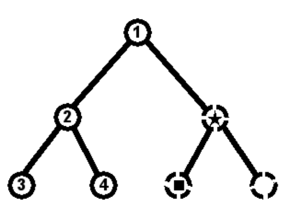

Fig. 3: Both DF and SH strategies visit nodes 1, 2, 3 and 4 in order. DF then visits the next unvisited node in the tree (star), whereas SH visits its nearby child (square) opportunistically. If (square) is interesting the rest of the children are visited. In case it is uninteresting the UAV visits the parent node (star) recursively.

Note that similar tree structures have been used before for e.g. probabilistic search with UAVs [18]. However we define a more general structure to be used in non-uniform coverage. The following section describes three strategies for traversing a coverage tree.

IV. TRAVERSAL STRATEGIES

As mentioned earlier, one way to cover the area is the lawnmower pattern at h_t in which all the leaves of the coverage tree are visited. However one can start at a higher altitude in the tree and visit the lower nodes if they happen to be interesting. In this case the range of the sensor should also be considered when choosing the highest altitude of the sensor, i.e. the robot should not go to a height at which the sensor can no longer distinguish the interesting child nodes from uninteresting ones. Let us call this highest level in the tree l_{max} which refers to a depth in the tree. We propose 3 approaches to traverse the coverage tree:

- **Breadth-First strategy (BF):** In this strategy the UAV starts at l_{max} and visits all the nodes in that level by performing a lawnmower pattern. Upon reaching each node, the corresponding child nodes are labeled as interesting or uninteresting. After the last node of l_{max} is visited, the UAV descends to the next depth in the tree and visits all nodes that were marked interesting in the previous traverse. Again all children are labeled. For path planning at each level we use a 2-approximation TSP tour. This process repeats until there is no node in the next depth of the tree.
- **Depth-First (DF):** In the DF strategy, similar to BF, the UAV starts performing the lawnmower trajectory at l_{max}, however, after visiting each node, the UAV descends to the next depth to visit all interesting children (if any). This behaviour is repeated recursively upon visiting each child node (see Fig 3).
- **Shortcut Heuristic (SH):** This strategy is the same as DF with one difference: Assume the UAV has visited a node and the next node to visit, n_{next}, is located at some height above the current node. According to DF the UAV flies directly to n_{next}. However in SH, the UAV visits the nearest child ($n_{nearest}$) of n_{next}. Now there are two possible situations: $n_{nearest}$ was interesting or it was uninteresting. In case it was interesting, we recursively visit all other unvisited children of n_{next} (see Fig 3). Note that we do not have to visit n_{next}. Alternatively if the child was uninteresting, the UAV visits the parent of $n_{nearest}$.

In the Breadth-First strategy, the whole area is covered initially with coarse resolution. The paths at higher levels are relatively short since the footprint of the sensor sweeps a larger area. Then the sub-regions that might be interesting in the next levels are determined and visited. It is worthwhile to note that this strategy provides a complete but coarse

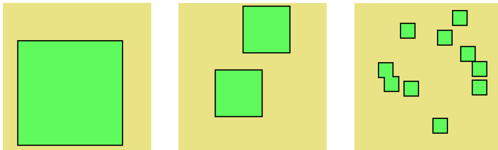

Fig. 4: Sample synthetic environments that are generated for the first experiment. The squares show interesting areas.

coverage of the whole area after a short time and gradually provides higher resolution coverage of the interesting regions. Previous work has assumed that such initial coarse coverage is available to be used as input to a planner [19], [20], [21], but our approach is a generalization of this 2-level method.

The Depth-First and the Shortcut Heuristic, on the other hand, provide the finest necessary coverage of a sub-region and then move on to the remaining parts of the area. Consequently, with DF and SH, the UAV never goes back to a previously visited region for higher resolution coverage. Intuitively, BF provides a complete but low resolution coverage as early as possible - an "anytime" strategy - while BF and SH provide the high resolution coverage in the shortest time. In many applications interesting regions form big patches, e.g. see Fig (2b). In such situations, the probability that one node is interesting provided a sibling node is interesting, will be high. This *locality* property is exploited by the Shortcut Heuristic. This causes the UAV to remain close to the lowest altitude in the interesting regions and unlike the Depth-First strategy, it eliminates unnecessary visits to the corresponding parent nodes.

In the next section we describe our experiments and compare the strategies with the standard lawnmower pattern as the baseline. Note that when the environment is interesting at all locations, the lawnmower strategy is optimal in path length.

V. EXPERIMENTS AND RESULTS

We performed two different sets of experiments in simulation. In the first set the interesting regions are generated at random from a known distribution. In the second set, the interesting regions are identified in images of natural terrains.

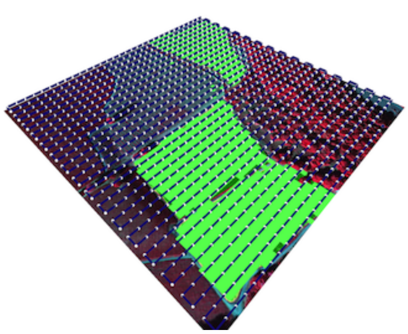

Fig. 6: Lawnmower-like pattern is used as the baseline approach. The area is covered with a uniform resolution.

A. Synthetic Environment

In these experiments an 128×128 m^2 area, free of obstacles, is considered (Figure 4). The simulated UAV is equipped with a sensor pointing downwards with $l(h) = h$, i.e., the field-of-view of 90°. In order to examine our methods in many environment configurations, two parameters are used to synthetically generate distributions of interesting regions: p, the percentage of the whole area that is interesting and c, the number of interesting segments. We assume that the interesting regions have the same size. For each pair of (p, c), $p \in \{10, .., 100\}$ and $c \in \{1, .., 10\}$, 10 random environments are generated and the four strategies (BF, DF, SH, lawnmower) are used to cover the area. For the traversal strategies we use $b = 2$, $l_{max} = 1$ and coverage tree with maximum depth of 5. For each strategy, the average of the total distance that the robot travelled is used as a performance metric. Figures 5 and 8 show the results of these experiments. In all experiments the variance was below 189 m (the error bars are very small and difficult to see in the figures).

The results indicate that if the interesting portion of the environment is less than 20% and forms few patches, (almost) all coverage tree strategies perform better than the lawnmower pattern. For example if 25% of the area is interesting then the shortcut heuristic takes a little more than half the length of the lawnmower pattern to cover the area. This means that the UAV could finish the coverage in half the time. As the number of interesting regions increases, the BF strategy becomes inefficient quickly. This is because the BF strategy visits all interesting nodes at the first tree level, then descends to the next level and revisits all nodes again and so forth. This also means uninteresting areas between interesting patches are revisited at higher resolutions. The revisiting and switching between patches of interest increases the travel cost quickly. In contrast DF and SH perform one high level (short distance) traverse and immediately exploit the available information to visit only interesting areas.

An area with large segments of interesting regions (i.e. small c) results in a high locality property in which SH outperforms the lawnmower and DF strategies. However, as the distribution of interesting regions varies from few large segments to a more uniform one, i.e. c becomes large, the locality property decreases and lawnmower becomes superior to both DF and SH. This is due to the fact that the lawnmower pattern never revisits a location.

Figure 8 shows the path length of each strategy for different percentages of interesting area in the environment. Only one single patch is used in these experiments. As the interesting patch grows, we expect that the Shortcut Heuristic keeps the UAV close to the ground to cover the interesting neighbourhood which yields a better performance compared to Depth-First. As the fraction of interesting area grows, the Shortcut Heuristic is eventually outperformed by the lawnmower which is optimal when 100% of the environment is interesting.

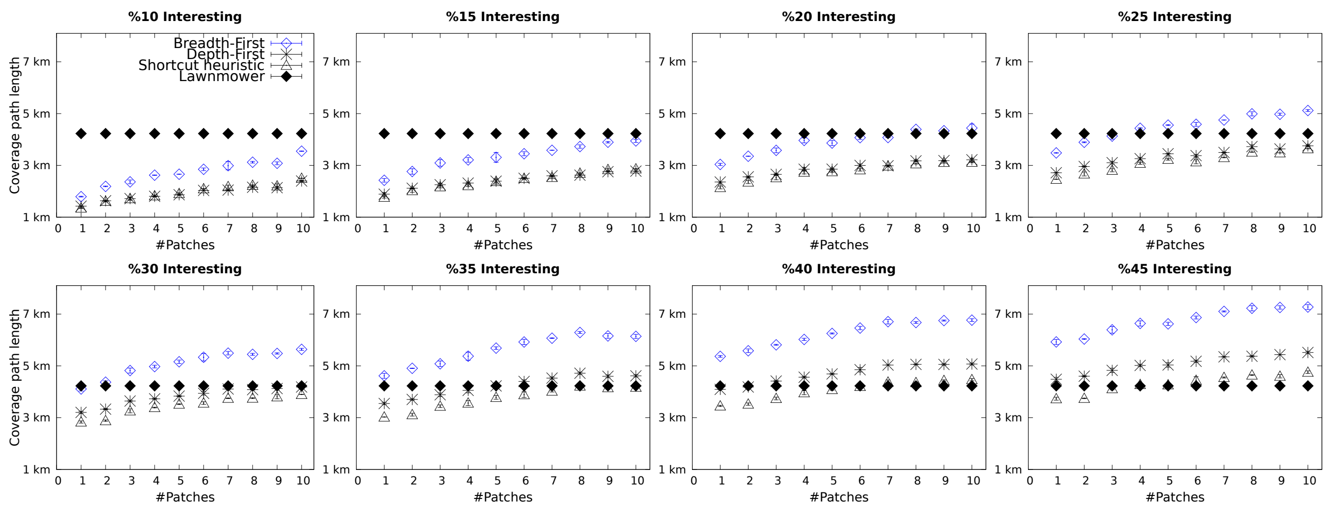

Fig. 5: Results of the simulations: each graph shows the average length of the coverage paths generated by each strategy for different distribution of interesting regions.

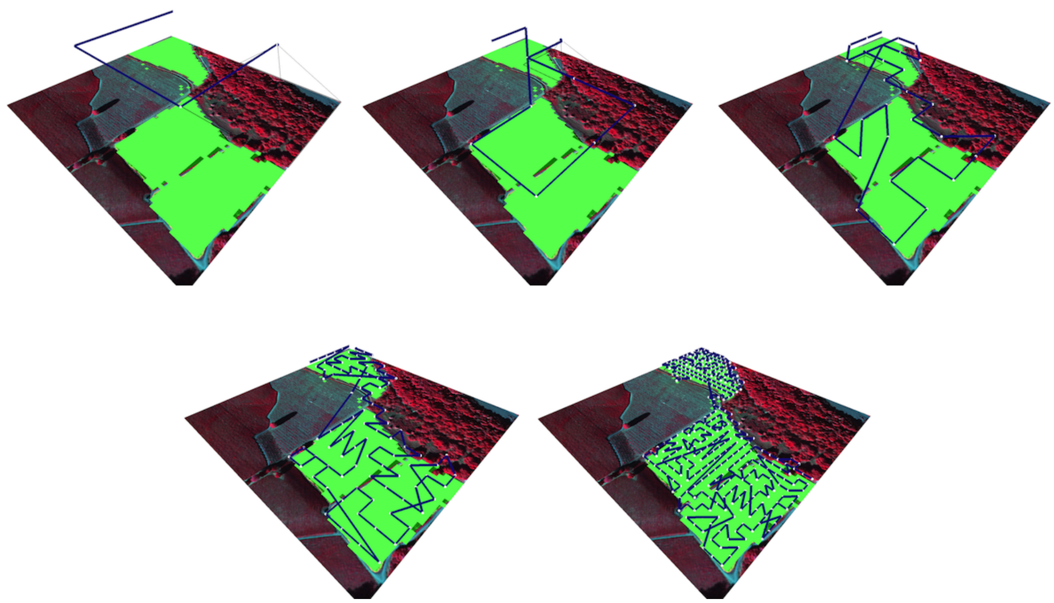

Fig. 7: Sample real environments are simulated in the second experiment. These figures show the coverage plan (generated by TSP 2-approximation) for each level of the tree in the BF strategy. At the first level with no prior interestingness information, coarse lawnmower pattern is performed. Light green (light grey) pixels are interesting.

B. Simulated Environment

We also used example images of natural environments for our simulated UAV. Figure 2 shows two sample environments. In Figure 2a the bushes and Figure 2b fields with a particular vegetation are considered interesting respectively. The interestingness detector checks if the color of a pixel falls in a certain color range. The area shown in Fig.2a represents an environment with a number of small interesting segments (Environment 1) whereas Figure 2b demonstrate an area with few large continuous interesting regions (Environment 2). We ran the same experiments as described in the previous section on these two environments, with results shown in Table I.

Strategy	Environment 1	Environment 2
Breadth-First	0.86	1.04
Depth-First	0.56	0.81
Shortcut Heuristic	0.59	0.69
Lawnmower pattern	1	1

TABLE I: Ratio of coverage paths lengths generated by the coverage tree strategies to the length of the lawnmower-like pattern.

The results indicate that the lawnmower pattern performs worse than the DF and SH in these maps. This is because the

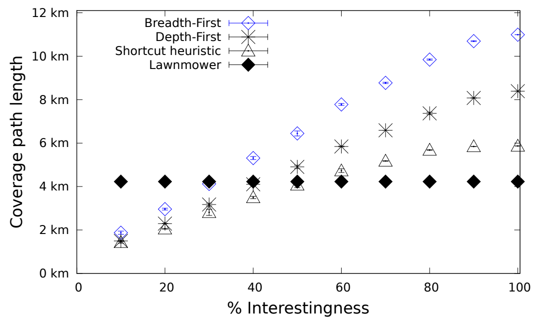

Fig. 8: Results of the experiments with a single patch.

majority of the area is uninteresting and the two mentioned strategies are able to behave accordingly. In the first environment (Fig. 2a), Depth-First performs a little better than the Shortcut Heuristic due to the low locality. However in the second map (Fig. 2b) the large interesting regions are covered more efficiently with the SH strategy compared to DF.

VI. CONCLUSION AND FUTURE WORK

Aerial coverage is a practical application of UAVs. Efficient coverage planning will increase the amount of useful area covered in a fixed flying time. In this paper we proposed methods that perform adaptive coverage according to the distribution of interesting regions in the environment. Simulation experiments show that our methods are effective in environments with certain types of interestingness distributions, namely sparse and patchy which are common in natural environments. We further showed that in some situations our two best methods outperformed the standard lawnmower approach by a factor of almost 1.7 overall.

In future we will perform experiments with real UAVs. We will also explore using lawnmower pattern for locally interesting patches. Furthermore, we will consider different energy consumption models for vertical/horizontal movements and compare the proposed strategies in terms of energy consumption. Also concrete metrics of interestingness should be developed to be used in real-world experiments. In this work, no time constraint was imposed and the assumption was that the UAV has enough battery life to finish the whole traverse. However in practice these vehicles (i.e. quadrotors) can only perform short flights and thus they are unable to perform complete coverage in large areas as discussed. Research on what would be a good strategy in these situations is very interesting.

ACKNOWLEDGMENT

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) through the NSERC Canadian Field Robotics Network (NCFRN).

REFERENCES

[1] A. Barrientos, J. Colorado, J. d. Cerro, A. Martinez, C. Rossi, D. Sanz, and J. Valente, “Aerial remote sensing in agriculture: A practical approach to area coverage and path planning for fleets of mini aerial robots,” *Journal of Field Robotics*, vol. 28, no. 5, pp. 667–689, 2011.

[2] L. Doitsidis, A. Renzaglia, S. Weiss, E. Kosmatopoulos, D. Scaramuzza, and R. Siegwart, “3d surveillance coverage using maps extracted by a monocular slam algorithm,” in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pp. 1661–1667, Sept 2011.

[3] T. Suzuki, Y. Amano, and T. Hashizume, “Generation of Large Mosaic Images for Vegetation Monitoring Using a Small Unmanned Aerial Vehicle,” *Journal of Robotics and Mechatronics*, vol. 22, no. 2, pp. 212–213, 2010.

[4] T.-S. Lee, J.-S. Choi, J.-H. Lee, and B.-H. Lee, “3-d terrain covering and map building algorithm for an auv,” in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pp. 4420–4425, Oct 2009.

[5] E. Galceran and M. Carreras, “A survey on coverage path planning for robotics,” *Robotics and Autonomous Systems*, vol. 61, pp. 1258–1276, Dec. 2013.

[6] E. Galceran and M. Carreras, “Efficient seabed coverage path planning for asvs and auvs,” in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 88–93, Oct 2012.

[7] J. Valente, D. Sanz, J. Del Cerro, A. Barrientos, and M. de Frutos, “Near-optimal coverage trajectories for image mosaicing using a mini quad-rotor over irregular-shaped fields,” *Precision Agriculture*, vol. 14, no. 1, pp. 115–132, 2013.

[8] J. Valente, A. Barrientos, J. Cerro, C. Rossi, J. Colorado, D. Sanz, and M. Garzn, “Multi-robot visual coverage path planning: Geometrical metamorphosis of the workspace through raster graphics based approaches,” in *Computational Science and Its Applications - ICCSA 2011*, vol. 6784 of *Lecture Notes in Computer Science*, pp. 58–73, Springer Berlin Heidelberg, 2011.

[9] I. Maza and A. Ollero, “Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms,” *Distributed Autonomous Robotic Systems 6*, pp. 221–230, 2007.

[10] S.-W. Moon and D. H.-C. Shim, “Study on path planning algorithms for unmanned agricultural helicopters in complex environment,” *International Journal of Aeronautical and Space Sicences*, vol. 10, no. 2, pp. 1–11, 2009.

[11] J. Cerro, A. Barrientos, D. Sanz, and J. Valente, “Aerial fleet in rhea project: A high vantage point contributions to robot 2013,” in *ROBOT2013: First Iberian Robotics Conference* (M. A. Armada, A. Sanfeliu, and M. Ferre, eds.), vol. 252, pp. 457–468, 2014.

[12] M. Johnson-Roberson, O. Pizarro, S. B. Williams, and I. Mahon, “Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys,” *Journal of Field Robotics*, vol. 27, no. 1, pp. 21–51, 2010.

[13] E. Galceran and M. Carreras, “Planning coverage paths on bathymetric maps for in-detail inspection of the ocean floor,” in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pp. 4159–4164, May 2013.

[14] E. Galceran and M. Carreras, “Coverage path planning for marine habitat mapping,” *2012 Oceans*, pp. 1–8, Oct. 2012.

[15] X. Zheng and S. Koenig, “Robot coverage of terrain with non-uniform traversability,” in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, pp. 3757–3764, IEEE, 2007.

[16] P. N. Atkar, D. C. Conner, A. Greenfield, H. Choset, and A. A. Rizzi, “Hierarchical segmentation of piecewise pseudoextruded surfaces for uniform coverage.,” *IEEE T. Automation Science and Engineering*, vol. 6, no. 1, pp. 107–120, 2009.

[17] S. Weiss, M. Achtelik, L. Kneip, D. Scaramuzza, and R. Siegwart, “Intuitive 3D Maps for MAV Terrain Exploration and Obstacle Avoidance,” *Journal of Intelligent & Robotic Systems*, vol. 61, pp. 473–493, Nov. 2010.

[18] S. Carpin, D. Burch, N. Basilico, T. H. Chung, and M. Klsch, “Variable resolution search with quadrotors: Theory and practice,” *Journal of Field Robotics*, vol. 30, no. 5, pp. 685–701, 2013.

[19] G. a. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, “Active planning for underwater inspection and the benefit of adaptivity,” *The International Journal of Robotics Research*, vol. 32, pp. 3–18, Nov. 2012.

[20] B. J. Englot, “Sampling-based coverage path planning for complex 3D structures,” 2012.

[21] C. Hoppe, A. Wendel, S. Zollmann, K. Pirker, A. Irschara, H. Bischof, and S. Kluckner, “Photogrammetric camera network design for micro aerial vehicles,” in *Computer Vision Winter Workshop (CVWW)*, vol. 8, pp. 1–3, 2012.